

Cross-correlation of CMB with large-scale structure: weak gravitational lensing

Christopher M. Hirata,^{1,*} Nikhil Padmanabhan,¹ Uroš Seljak,¹ David Schlegel,² and Jonathan Brinkmann³

¹*Department of Physics, Jadwin Hall, Princeton University, Princeton, New Jersey 08544, USA*

²*Princeton University Observatory, Princeton University, Princeton, New Jersey 08544, USA*

³*Apache Point Observatory, 2001 Apache Point Road, Sunspot, New Mexico 88349-0059, USA*

(Dated: August 30, 2004)

We present the results of a search for gravitational lensing of the cosmic microwave background (CMB) in cross-correlation with the projected density of luminous red galaxies (LRGs). The CMB lensing reconstruction is performed using the first year of *Wilkinson Microwave Anisotropy Probe* (WMAP) data, and the galaxy maps are obtained using the Sloan Digital Sky Survey (SDSS) imaging data. We find no detection of lensing; our constraint on the galaxy bias derived from the galaxy-convergence cross-spectrum is $b_g = 1.81 \pm 1.92$ (1σ , statistical), as compared to the expected result of $b_g \sim 1.8$ for this sample. We discuss possible instrument-related systematic errors and show that the Galactic foregrounds are not important. We do not find any evidence for point source or thermal Sunyaev-Zel'dovich effect contamination.

PACS numbers: 98.80.Es, 98.62.Sb, 98.62.Py

I. INTRODUCTION

The *Wilkinson Microwave Anisotropy Probe* (WMAP) [108] satellite has provided a wealth of information about the universe through its high-resolution, multi-frequency, all-sky maps of the cosmic microwave background (CMB) [1]. While the WMAP power spectrum [2] and temperature-polarization cross-spectrum [3] are useful for probing the high-redshift universe (reionization and earlier epochs) [4, 5, 6, 7], the WMAP maps also provide an opportunity to study the low-redshift universe through secondary CMB anisotropies. While the effect of secondary anisotropies on the angular scales probed by WMAP ($l \lesssim 700$) is small compared to the primordial temperature fluctuations, the signal-to-noise ratio can be boosted by cross-correlating with tracers of the large scale structure (LSS) at low redshifts. Since the WMAP data release, several authors have used various tracers of LSS to measure the integrated Sachs-Wolfe (ISW) effect, the thermal Sunyaev-Zel'dovich (tSZ) effect, and microwave point sources [8, 9, 10, 11, 12, 13, 14]. The Sloan Digital Sky Survey (SDSS) [109] is an excellent candidate for these cross-correlation studies due to the large solid angle covered at moderate depth.

Another secondary anisotropy, which has not yet been investigated observationally, is weak lensing of the CMB by intervening large-scale structure. Weak lensing has attracted much attention recently as a means of directly measuring the matter power spectrum at low redshifts (e.g. [15]). The traditional approach is to use distant galaxies as the “sources” that are lensed to measure e.g. the matter power spectrum (e.g. [16, 17, 18, 19, 20, 21, 22, 23]) or the galaxy-matter cross-correlation (e.g. [24, 25, 26, 27, 28, 29]). However weak lensing of CMB offers an alternative method, free of intrinsic

alignments, uncertainties in the source redshift distribution, and selection biases (since the CMB is a random field). Potential applications of CMB lensing described in the literature include precision measurement of cosmological parameters [30, 31, 32, 33] and separation of the lensing contribution to the CMB B -mode polarization from primordial vector [34] and tensor perturbations [35, 36, 37]. While these applications are in the future, the WMAP data for the first time allows a search for weak lensing of the CMB in correlation with large-scale structure. This paper presents the results of such a search; our objective here is not precision cosmology, but rather to detect and characterize any systematic effects that contaminate the lensing signal at the level of the current data. This step is a prerequisite to future investigations that will demand tighter control of systematics.

In this paper, we perform cross-correlation analysis between the CMB weak lensing field derived from WMAP and a photometrically selected sample of luminous red galaxies (LRGs) in the SDSS at redshifts $0.2 < z < 0.7$. The photometric LRGs are well-suited for cross-correlation studies because of their high intrinsic luminosity (compared to normal galaxies), which allows them to be observed at large distances; their high number density, which suppresses shot noise in the maps; and their uniform colors which allow for accurate photometric redshifts and hence determination of the redshift distribution. We use the measured cross-spectrum between the lensing field and the projected galaxy density to estimate the LRG bias b_g . At the present stage, we are using the bias as a proxy for the strength of the cross-correlation signal, just as has been done in recent analyses of the ISW effect [8, 10, 12, 13, 14]; we are not yet trying to use the bias in cosmological parameter estimation, although this is a possible future application of the methodology. We do not have a detection of a cross-correlation, and hence our measured bias $b_g = 1.81 \pm 1.92$ is consistent with zero.

This paper is organized as follows. The most impor-

*Electronic address: chirata@princeton.edu

tant aspects (for this analysis) of the *WMAP* and SDSS data sets, and the construction of the LRG catalog, are described in Sec. II. The theory of CMB lensing and reconstruction methodology are explained in Sec. III. The cross-correlation methodology and simulations are covered in Sec. IV, and the results are presented in Sec. V. We investigate possible systematic errors in Sec. VI, and conclude in Sec. VII. Appendix A describes the spherical harmonic transform algorithms and associated conventions used in this paper, and Appendix B describes the algorithm used for the C^{-1} operations that arise in our analysis.

II. DATA

A. CMB temperature from *WMAP*

The *WMAP* mission [38] is designed to produce all-sky maps of the CMB at multipoles up to $l \sim$ several hundred. This analysis uses the first public release of *WMAP* data, consisting of one year of observations from the Sun-Earth L2 Lagrange point. *WMAP* carries ten differencing assemblies (DAs), each of which measures the difference in intensity of the CMB at two points on the sky; a CMB map is built up from these temperature differences as the satellite rotates. (*WMAP* has polarization sensitivity but this is not used in the present analysis.) The DAs are designated K1, Ka1, Q1, Q2, V1, V2, W1, W2, W3, and W4; the letters indicate the frequency band to which a particular DA is sensitive [1, 39] (the K, Ka, Q, V, and W bands correspond to central frequencies of 23, 33, 41, 61, and 94 GHz, respectively). The *WMAP* team has pixelized the data from each DA in the HEALPix [110] pixelization system at resolution 9 [1, 40]. This system has 3,145,728 pixels, each of solid angle 47.2 sq. arcmin. These maps are not beam-deconvolved; this, combined with the *WMAP* scan strategy, results in nearly uncorrelated Gaussian uncertainties on the temperature in each pixel.

In this paper, we only use the three high-frequency microwave bands (Q, V, and W) because the K and Ka bands are very heavily contaminated by galactic foregrounds and have poor resolution. (The foreground emission is not a Gaussian field and cannot be reliably simulated, so in cases where it dominates over CMB anisotropy and instrument noise, we cannot compute reliable error bars on the cross-correlation.) For the galaxy-lensing correlation, we have used the sky maps produced by the eight high-frequency DAs. The variances of the temperature measurements are obtained from the effective number of observations N_{obs} .

Note that the *WMAP* “internal linear combination” (ILC) map [1] cannot be used for lensing studies because of its degraded resolution (1 degree full width half maximum, FWHM), which eliminates the multipoles $l \sim 350$ of greatest importance for the lensing analysis. An ILC-based lensing analysis would also suffer from practical

issues, namely the loss of frequency-dependent information (useful as a test of foregrounds), the inability to separate cross-correlations between different DAs from auto-correlations (useful to avoid the need for noise bias subtraction), and the complicated inter-pixel noise correlations (due to the smoothing used to create the map and the varying weights of the different frequencies). The foreground-cleaned map of Ref. [41] recovers the full *WMAP* resolution, but the practical difficulties (for the purpose of lensing reconstruction) associated with ILC still apply. We have not used either of these maps in this paper.

B. LRG density from SDSS

The Sloan Digital Sky Survey (SDSS) [42] is an ongoing survey to image approximately π steradians of the sky, and follow up approximately one million of the detected objects spectroscopically [43, 44]. The imaging is carried out by drift-scanning the sky in photometric conditions [45], in five bands (*ugriz*) [46, 47] using a specially designed wide-field camera [48]. These imaging data are the source of the LSS sample that we use in this paper. In addition, objects are targeted for spectroscopy using these data [49] and are observed with a 640-fiber spectrograph on the same telescope. All of these data are processed by completely automated pipelines that detect and measure photometric properties of objects, and astrometrically calibrate the data [50, 51]. The SDSS is well underway, and has had three major data releases [52, 53, 54, 55]; this paper uses all data observed through Fall 2003 (296,872 HEALPix resolution 9 pixels, or 3,893 square degrees).

The SDSS detects many extragalactic objects that could, in principle, be used for cross-correlation with secondary anisotropies [56]. The usefulness of LRGs as a cosmological probe has been appreciated by a number of authors [57, 58]. These are typically the most luminous galaxies in the universe, and therefore probe cosmologically interesting volumes. In addition, these galaxies are generically old stellar systems and have extremely uniform spectral energy distributions (SEDs), characterized only by a strong discontinuity at 4000 Å. The combination of these two characteristics make them an ideal candidate for photometric redshift algorithms, with redshift accuracies of $\sigma_z \sim 0.03$ [59]. We briefly outline the construction of the photometric LRG sample used in this paper below, and defer a detailed discussion of the selection criteria and properties of the sample to a later paper [60].

Our selection criteria are derived from those described in Ref. [58]. However, since we are working with a photometric sample, we are able to relax the apparent luminosity constraints imposed there to ensure good throughput on the SDSS spectrographs. We select LRGs by choosing galaxies that both have colors consistent with an old stellar population, as well as absolute luminosities greater

than a chosen threshold. The first criterion is simple to implement since the uniform SEDs of LRGs imply that they lie on an extremely tight locus in the space of galaxy colors; we simply select all galaxies that lie close to that locus. More specifically, we can define three (not independent) colors that describe this locus,

$$\begin{aligned} c_{\perp} &\equiv (r - i) - 0.25(g - r) - 0.18, \\ d_{\perp} &\equiv (r - i) - 0.125(g - r), \\ c_{||} &\equiv 0.7(g - r) + 1.2(r - i - 0.18), \end{aligned} \quad (1)$$

where g , r , and i are the SDSS model magnitudes [52] in the g , r and i bands (centered at 469, 617, and 748 nm respectively). We now make the following color selections,

$$\begin{aligned} \text{Cut I: } & |c_{\perp}| < 0.2; \\ \text{Cut II: } & d_{\perp} > 0.55, \quad g - r > 1.4. \end{aligned} \quad (2)$$

Making two cuts (Cut I and Cut II) is convenient since the LRG color locus changes direction sharply as the 4000 Å break redshifts from the g to the r band; this division divides the sample into a low redshift (Cut I, $z < 0.4$) and high redshift (Cut II, $z > 0.4$) sample.

In order to implement the absolute magnitude cut, we follow [58] and impose a cut in the galaxy color-magnitude space. The specific cuts we use are

$$\begin{aligned} \text{Cut I: } & r_{\text{Petro}} < 13.6 + \frac{c_{||}}{0.3}, \quad r_{\text{Petro}} < 19.7, \\ \text{Cut II: } & i < 18.3 + 2d_{\perp}, \quad i < 20, \end{aligned} \quad (3)$$

where r_{Petro} is the SDSS r band Petrosian magnitude [52]. Finally, we reject all objects that resemble the point-spread function of the telescope, or if they have colors inconsistent with normal galaxies; these cuts attempt to remove interloping stars.

Applying these selection criteria to the ~ 5500 degrees of photometric SDSS imaging in the Galactic North yields a catalog of approximately 900,000 galaxies. Applying the single template fitting photometric redshift algorithm of [59], we restrict this catalog to galaxies with $0.2 < z_{\text{photo}} < 0.6$, leaving us with $\sim 650,000$ galaxies. We use the regularized inversion method of Ref. [59] as well as the photometric redshift error distribution presented there, to estimate the true redshift distribution of the sample. The results, comparing the photometric and true redshift distributions are shown in Fig. 1. Finally, this catalog is pixelized as a number overdensity, $g = \delta n / \bar{n}$, onto a HEALPix pixelization of the sphere, with 3,145,728 pixels. We also mask regions around stars from the Tycho astrometric catalog [61], as the photometric catalogs are incomplete near bright stars. The final catalog covers a solid angle of 3,893 square degrees (296,872 HEALPix resolution 9 pixels) and contains 503,944 galaxies at a mean density of 1.70 galaxies per pixel.

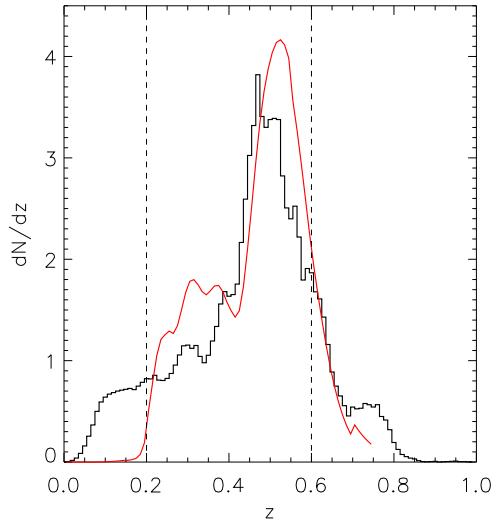


FIG. 1: The LRG redshift distribution. The black histogram shows the photo- z distribution, the red curve is the true redshift distribution estimated by regularized deconvolution of the photo- z errors.

III. LENSING OF CMB

A. Definitions

Gravitational lensing re-maps the primordial CMB anisotropy \tilde{T} into a lensed temperature T according to

$$T(\hat{\mathbf{n}}) = \tilde{T}(\hat{\mathbf{n}} + \mathbf{d}(\hat{\mathbf{n}})), \quad (4)$$

where the 2-vector \mathbf{d} is the deflection angle of null geodesics. To first order in the metric perturbations, \mathbf{d} can be expressed as the gradient of a scalar lensing potential, $\mathbf{d} = \nabla\Phi$, where ∇ is the derivative on the unit (celestial) sphere. We may also define the convergence $\kappa = -\frac{1}{2}\nabla \cdot \mathbf{d}$. Assuming the primordial CMB is statistically isotropic with some power spectrum C_l , it can be shown [62, 63] that the multipole moments of the lensed temperature field have covariance

$$\begin{aligned} \langle T_{l_1 m_1}^* T_{l_2 m_2} \rangle &= C_l \delta_{l_1 l_2} \delta_{m_1 m_2} + \sum_{LM} (-1)^{m_2} \mathcal{J}_{L l_1 l_2} \\ &\times \begin{pmatrix} l_1 & l_2 & L \\ -m_1 & m_2 & -M \end{pmatrix} \kappa_{LM}, \end{aligned} \quad (5)$$

where we have introduced the Wigner $3j$ symbol, and the coupling coefficient is

$$\begin{aligned} \mathcal{J}_{L l_1 l_2} &= \frac{2}{L(L+1)} \sqrt{\frac{(2L+1)(2l_1+1)(2l_2+1)}{16\pi}} \\ &\times \left\{ [L(L+1) + l_1(l_1+1) - l_2(l_2+1)]C_{l_1} \right. \\ &+ [L(L+1) - l_1(l_1+1) + l_2(l_2+1)]C_{l_2} \left. \right\} \\ &\times \begin{pmatrix} l_1 & l_2 & L \\ 0 & 0 & 0 \end{pmatrix}. \end{aligned} \quad (6)$$

The *WMAP* satellite does not directly measure T , but rather a beam-convolved temperature:

$$\hat{T}^\alpha(\hat{\mathbf{n}}) = \int B^\alpha(\hat{\mathbf{n}}, \hat{\mathbf{n}}') T(\hat{\mathbf{n}}') d^2 \hat{\mathbf{n}}' + \text{noise}. \quad (7)$$

(Here $\alpha = \text{K1..W4}$ are the differencing assemblies.) In most of this analysis we have approximated the beam by the *WMAP* circularized beam transfer function [64]. The temperature multipole moments recovered assuming a circular beam are

$$\hat{T}_{lm}^\alpha = \frac{\dot{T}_{lm}^\alpha}{B_l^\alpha}, \quad (8)$$

where B_l^α are the beam transfer functions. If the beam is truly circular, Eq. (8) returns an unbiased estimator of the beam-deconvolved CMB temperature; in Sec. VI C, we will consider the effect of the *WMAP* beam ellipticity on lensing estimation. Note further that \hat{T}_{lm}^α is only well-determined up to some maximum multipole l because the B_l^α drop to zero at high l .

B. Theoretical predictions for lensing

In this paper, we aim to measure the galaxy-convergence cross-correlation, $C_l^{g\kappa}$, where $g = \delta n/\bar{n}$ is the projected fractional overdensity of galaxies; this section briefly presents the theoretical prediction from the Λ CDM cosmology. In a spatially flat Friedmann-Robertson-Walker universe described by general relativity, the convergence is given in terms of the fractional density perturbation δ by:

$$\kappa(\hat{\mathbf{n}}) = 4\pi G_N \bar{\rho}_0 \int \frac{\chi(\chi_{CMB} - \chi)}{\chi_{CMB}} (1+z) \delta(\chi, \hat{\mathbf{n}}) d\chi, \quad (9)$$

where χ is the comoving radial distance, z is the redshift observed at radial distance χ , $\bar{\rho}_0$ is the present-day mean density of the universe, and χ_{CMB} is the comoving distance to the CMB. The galaxy overdensity does not come from a “clean” theoretical prediction, but on large scales it can be approximated by

$$g(\hat{\mathbf{n}}) = \frac{\int b_g \mathcal{N}(\chi) \delta(\chi, \hat{\mathbf{n}}) d\chi}{\int \mathcal{N}(\chi) d\chi}, \quad (10)$$

where $\mathcal{N}(\chi)$ is the distribution in comoving distance and b_g is the galaxy bias. (The SDSS LRG sample is at low redshift $z \leq 0.7$ and does not have a steep luminosity function at the faint end of our sample, so we neglect the magnification bias.) For $l \gg 1$ and smooth power spectra for matter and galaxies, this may be approximated by the Limber integral:

$$C_l^{g\kappa} = 4\pi G_N \bar{\rho}_0 \times \frac{\int b_g \mathcal{N}(\chi) \left(\frac{1}{\chi} - \frac{1}{\chi_{CMB}} \right) (1+z) P_\delta(\frac{l}{\chi}) d\chi}{\int \mathcal{N}(\chi) d\chi}, \quad (11)$$

where the matter power spectrum P_δ is evaluated at comoving wavenumber $k = l/\chi$ and at the redshift corresponding to conformal time $\eta_0 - \chi$. It is obtained using the transfer functions from CMBFAST [65] and the best-fit 6-parameter flat Λ CDM cosmological model from *WMAP* and SDSS data from Ref. [66] ($\Omega_b h^2 = 0.0232$; $\Omega_m h^2 = 0.1454$; $h = 0.695$; $\tau = 0.124$; $\sigma_8 = 0.917$; $n_s = 0.977$). We have found that varying each of these parameters over their 1σ uncertainty ranges gives a $\pm 18\%$ effect for σ_8 (for which $C_l^{g\kappa}$ scales as $\propto \sigma_8^2$) and $\pm < 14\%$ effect for the other parameters. Since the overall significance of the detection is only 0.9σ , this dependence of the template $C_l^{g\kappa}$ on cosmological parameters will be neglected here.

The lensing signal is on large scales and so we have not used a nonlinear mapping. The peak of the LRG redshift distribution is at $z \sim 0.5$, corresponding to a comoving angular diameter distance $\sim 1.3h^{-1}$ Gpc, in which case the smallest angular scales we use ($l = 300$) correspond to $k = 0.23h \text{ Mpc}^{-1}$ and $\Delta^2(k) = 0.7$. The nonlinear evolution at this scale according to the Peacock-Dodds formula [67] is a 10% correction to the matter power spectrum and is thus much smaller than the error bars presented in this paper (although it is not obvious what this implies about the galaxy-matter cross-spectrum, which is the quantity that should appear in Eq. 11). Future applications of CMB lensing in precision cosmology will of course require accurate treatment of the nonlinear evolution.

In this paper, we will assume that galaxy bias b_g is constant so that it may be pulled out of the integrals in Eqs. (10) and (11). If the bias varies with redshift (as suggested by e.g. [68]), then the best-fit value of b_g will be some weighted average of b_g over the redshift distribution; this need not be the same weighted average that one observes from the auto-power spectrum, since the latter is weighted differently. Computation of the auto-power in photometric redshift slices [60] suggests that over the redshift range $0.2 \leq z \leq 0.6$ the bias varies from 1.7–1.9; this variation can safely be neglected given our current statistical errors. Note however that a detection of $b_g \neq 0$ via a fit to Eq. (11) assuming constant b_g would rule out $C_l^{g\kappa} = 0$ and hence would be robust evidence for a galaxy-convergence correlation, regardless of the redshift dependence of the bias.

C. Lensing reconstruction

We construct lensing deflection maps using quadratic reconstruction methods [30, 63, 69, 70, 71, 72, 73, 74, 75], which have been shown to be near-optimal for lensing studies of the CMB temperature on large ($l < 3500$) scales [76]. Non-quadratic methods may be superior if CMB polarization is used [77], or on very small scales [78, 79, 80]; these cases are not of interest for *WMAP*, since the sensitivity is insufficient to map the CMB polarization and arcminute scales are unresolved. Quadratic

estimation takes advantage of the cross-coupling of different multipoles induced by gravitational lensing, namely the $O(\kappa_{LM})$ term in Eq. (5). The maximum signal-to-noise statistic for CMB weak lensing is the divergence of the temperature-weighted gradient. We construct, for each pair of differencing assemblies α and β , the temperature-weighted gradient vector field $\tilde{\mathbf{G}}^{\alpha\beta}$:

$$\begin{aligned}\tilde{\mathbf{G}}^{\alpha\beta}(\hat{\mathbf{n}}) = \frac{1}{2} & \left([W\hat{T}^\alpha](\hat{\mathbf{n}}) \nabla [CWT\hat{T}^\beta](\hat{\mathbf{n}}) \right. \\ & \left. + [W\hat{T}^\beta](\hat{\mathbf{n}}) \nabla [CWT\hat{T}^\alpha](\hat{\mathbf{n}}) \right),\end{aligned}\quad (12)$$

where $W\hat{T}$ and $CWT\hat{T}$ are defined by the convolution relations $[W\hat{T}^\alpha]_{lm} = W_l \hat{T}_{lm}^\alpha$ and $[CWT\hat{T}^\alpha]_{lm} = C_l W_l \hat{T}_{lm}^\alpha$. Note that Eq. (12) is exactly the same as the $\tilde{\mathbf{G}}$ statistic of Ref. [73] except that we have multiple differencing assemblies, and we have left open the choice of the weight W_l . While $W_l = (C_l + C_l^{noise})^{-1}$ is statistically optimal, there are also practical considerations that affect this choice. Specifically, it is desirable that the $W\hat{T}$ and $CWT\hat{T}$ convolutions are almost-local functions of the CMB temperature (to minimize leakage from the Galactic Plane), and that the same W_l be used for all differencing assemblies (so that any frequency dependence of our results can be attributed to foregrounds or noise, rather than merely a change in which primary CMB modes we are studying). We choose the following weight function:

$$W_l = \left[C_l + (0.03461 \mu\text{K}^2) e^{l(l+1)/300^2} \right]^{-1}, \quad (13)$$

which clearly has the optimal C_l^{-1} dependence in the high signal-to-noise regime. Note that in the range of l we use ($l \leq 800$), the W_l drop to zero with increasing l faster than the Q-band beam transfer functions. Hence the computation of $W\hat{T}^{Q1,Q2}$ are stable even though $\hat{T}^{Q1,Q2}$ are beam-deconvolved. Due to their narrower beams, this stability also applies to the V and W-band DAs. The K1 and Ka1 DAs have wider beams and hence lensing reconstruction using the weight Eq. (13) is unstable for these DAs. We set $W_0 = W_1 = 0$ to reject the monopole (not observed by *WMAP*) and dipole signals. The power spectrum C_l used for the lensing reconstruction is *WMAP* best-fit ΛCDM model with scalar spectral running α_s [4] to the CMB data (*WMAP*+ACBAR [81]+CBI [82]). Errors in the C_l used in the lensing reconstruction cannot produce a spurious galaxy-temperature correlation because they result only in a calibration error in the lensing estimator. Furthermore, *WMAP* has determined the C_l to within several percent (except at the low multipoles, which give a subdominant contribution to both $W\hat{T}$ and $\nabla CWT\hat{T}$), whereas the lensing cross-correlation signal is only present at the 1σ level, this error is not important for the present analysis.

In the reconstruction method of Ref. [73], a filtered divergence of $\tilde{\mathbf{G}}$ is taken to extract the lensing field. We avoid this step because it is highly nonlocal, and hence

can smear Galactic plane contamination into the regions of sky used for lensing analysis. In principle, we would prefer to directly cross-correlate $\tilde{\mathbf{G}}$ with the LRG map, but this too is difficult because the noise power spectrum of $\tilde{\mathbf{G}}$ is extremely blue. We compromise by computing \mathbf{v} , a Gaussian-filtered version of $\tilde{\mathbf{G}}$:

$$v_{lm}^{(\parallel, \perp)} = e^{-l(l+1)\sigma_0^2/2} \tilde{G}_{lm}^{(\parallel, \perp)}. \quad (14)$$

Here \parallel and \perp represent the longitudinal (vector) and transverse (axial) multipoles, which are the vector analogues of the tensor E and B multipoles. The Gaussian filter eliminates the troublesome high- l power present in $\tilde{\mathbf{G}}(\hat{\mathbf{n}})$ and makes \mathbf{v} suitable for cross-correlation studies. We have chosen a width $\sigma_0 = 0.01$ radians (34 arcmin).

The vector field \mathbf{v} can be written in terms of the temperatures \hat{T} directly in harmonic space. The longitudinal components are given by

$$\begin{aligned}v_{lm}^{\alpha\beta(\parallel)} = & (-1)^m \sum_{l'l''} \mathcal{K}_{ll'l''} \sum_{m'm''} \begin{pmatrix} l & l' & l'' \\ -m & m' & m'' \end{pmatrix} \\ & \times \frac{\hat{T}_{l'm'}^\alpha \hat{T}_{l''m''}^\beta + \hat{T}_{l'm'}^\beta \hat{T}_{l''m''}^\alpha}{2},\end{aligned}\quad (15)$$

where we have defined:

$$\begin{aligned}\mathcal{K}_{ll'l''} = & \begin{pmatrix} l & l' & l'' \\ 0 & 0 & 0 \end{pmatrix} \sqrt{\frac{(2l+1)(2l'+1)(2l''+1)}{16\pi l(l+1)}} \\ & \times [l(l+1) - l'(l'+1) + l''(l''+1)] C_{l''} W_{l'} W_{l''} \\ & \times e^{-l(l+1)\sigma_0^2/2}.\end{aligned}\quad (16)$$

We will not need the formula for the transverse components. While we perform the reconstruction (Eqs. 12 and 14) in real space, the harmonic-space relation (Eq. 15) is useful for computing the response of the estimator, and for estimating foreground contamination and beam effects. In particular, from the orthonormality relations for Wigner 3j-symbols, we have

$$\langle v_{lm}^{\alpha\beta(\parallel)} \rangle = \sum_{l'l''} \frac{\mathcal{J}_{ll'l''} \mathcal{K}_{ll'l''}}{2l+1} \kappa_{lm} \equiv R_l \kappa_{lm}, \quad (17)$$

which defines the calibration of \mathbf{v} as an estimator of the lensing field. The response factor R_l is shown in Fig. 2; we have verified this response factor in simulations (Sec. IV C).

There are 36 pairs of differencing assemblies that could be used to produce estimated lensing maps $\mathbf{v}_{lm}^{\alpha\beta}$: the 8 “autocorrelations” ($\alpha = \beta$) and 28 “cross-correlations” ($\alpha \neq \beta$). Note that instrument noise that is non-uniform across the sky can produce a bias in the “autocorrelation”-derived lensing map $\mathbf{v}^{\alpha\alpha}$. In principle the noise bias could be estimated and subtracted, just as can be done for the power spectrum. However, this is dangerous if the noise properties are not very well modeled. Since the *WMAP* noise is in fact strongly variable across the sky we only use the “cross-correlation” $\alpha \neq \beta$ maps.

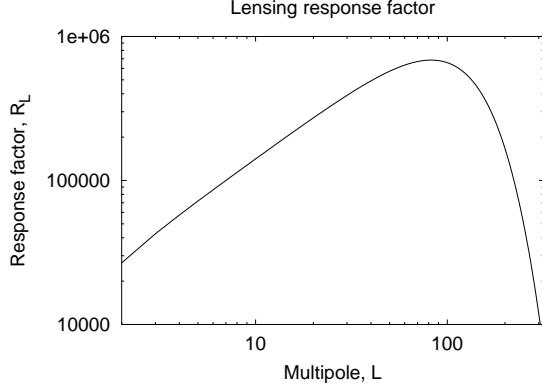


FIG. 2: The response factor R_l of Eq. (17) satisfying $\langle v_{lm}^{\alpha\beta}(\parallel) \rangle = R_l \kappa_{lm}$.

One problem that we find with this method is that the \mathbf{v} field contains “ghosts” caused by the Galactic plane (where small-scale temperature fluctuations of several millikelvin or more can occur due to Galactic emission). We solve this problem by setting $T = 0$ within the WMAP Kp4 [83] Galactic plane mask. We have verified that using the Kp2 mask instead produces only small changes to the results.

The weight functions W_l and $C_l W_l$ are shown in Fig. 3. We also show the real-space weights, given by

$$W(\theta) = \sum_l \frac{2l+1}{4\pi} W_l P_l(\cos \theta), \quad (18)$$

and similarly for $[CW](\theta)$.

D. Frequency-averaged lensing maps

The methodology outlined in Sec. III C allows us to construct 28 lensing maps $\mathbf{v}^{\alpha\beta}$ corresponding to the 28 pairs of differencing assemblies. For this analysis, we need to produce an “averaged” lensing map $\mathbf{v}^{(TT)}$ based on a minimum-variance linear combination of the 28 DA-pair maps. The averaged lensing map is determined by the weights $a_{\alpha\beta}^{(TT)}$:

$$\mathbf{v}^{(TT)} = \sum_{\alpha\beta} a_{\alpha\beta}^{(TT)} \mathbf{v}^{\alpha\beta}. \quad (19)$$

We select these weights to minimize the amount of power in $\mathbf{v}^{(TT)}$, subject to the restriction $\sum a_{\alpha\beta}^{(TT)} = 1$; this is done by minimizing the total vector power in \mathbf{v} between multipoles 50 and 125: $\mathcal{P} = \sum_{l=50}^{125} \sum_{m=-l}^l |v_{lm}^{(TT)}|^2$, which is a quadratic function of the weights $a_{\alpha\beta}^{(TT)}$. The optimal weights $a_{\alpha\beta}^{(TT)}$ are complicated to establish analytically since the maps $\mathbf{v}^{\alpha\beta}$ are highly correlated. We have therefore minimized \mathcal{P} using a *simulated* lensing map (see Sec. IV C). Using a simulated map rather than

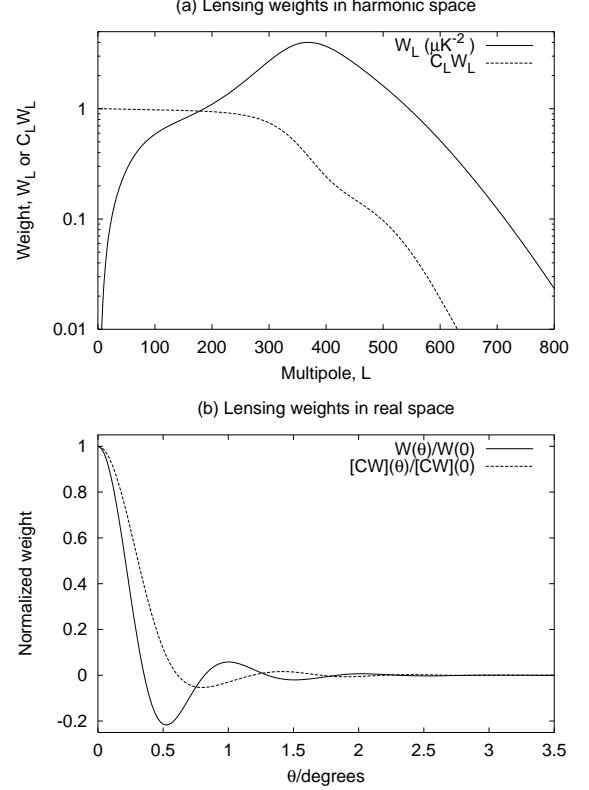


FIG. 3: (a) The weight functions W_l and $C_l W_l$. (b) The same weight functions in real-space (Eq. 18); $W\hat{T}$ and $CW\hat{T}$ are obtained by convolving the (beam-deconvolved) temperature \hat{T} with these kernels.

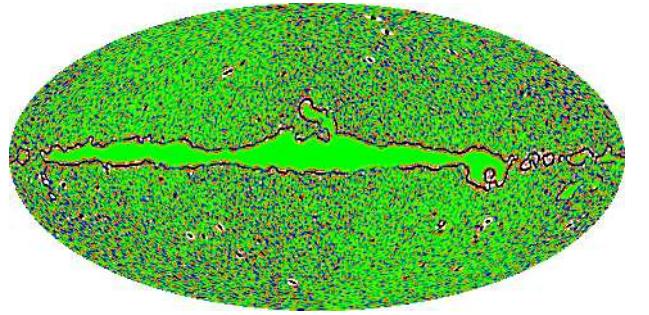


FIG. 4: The divergence of the lensing vector field map, $\nabla \cdot \mathbf{v}^{(TT)}$, smoothed with a 30 arcmin FWHM Gaussian, and displayed in Galactic Molleweide projection. Note the prominent artifacts surrounding the Galactic plane cut and the point sources (which are removed by the Kp05 \cap S10 \setminus ps2 cut).

the real data avoids the undesirable possibility of the weights being statistically correlated with the data. We also fix $a_{Q1,Q2}^{(TT)} = 0$ because the $\mathbf{v}^{Q1,Q2}$ map would be the most heavily contaminated by point sources. The weights so obtained are shown in Table I. A map of $\nabla \cdot \mathbf{v}^{(TT)}$, smoothed to 30 arcmin resolution (Gaussian FWHM) is shown in Fig. 4.

Also, to study foreground effects on lensing estima-

TABLE I: The weights $a_{\alpha\beta}$ used for the overall-averaged lensing map $a_{\alpha\beta}^{(TT)}$, and the six “individual frequency” maps $a_{\alpha\beta}^{(QQ)}, \dots, a_{\alpha\beta}^{(WW)}$.

DA pair $(\alpha\beta)$	$a_{\alpha\beta}^{(TT)}$	$a_{\alpha\beta}^{(\nu_1\nu_2)}$
Q1,Q2	0.000000	1.000000
Q1,V1	0.063759	0.220605
Q1,V2	0.075370	0.267859
Q2,V1	0.061538	0.227341
Q2,V2	0.098884	0.284196
Q1,W1	0.031583	0.167730
Q1,W2	0.026896	0.113943
Q1,W3	0.018391	0.094086
Q1,W4	0.028571	0.130429
Q2,W1	0.023816	0.158925
Q2,W2	0.017153	0.107379
Q2,W3	0.014494	0.097062
Q2,W4	0.027164	0.130444
V1,V2	0.106330	1.000000
V1,W1	0.053278	0.150880
V1,W2	0.029381	0.094494
V1,W3	0.029664	0.087425
V1,W4	0.035463	0.107193
V2,W1	0.048598	0.168035
V2,W2	0.049979	0.137365
V2,W3	0.034621	0.115438
V2,W4	0.046029	0.139169
W1,W2	0.016453	0.199003
W1,W3	0.014448	0.170408
W1,W4	0.019908	0.232958
W2,W3	0.014514	0.129381
W2,W4	0.003365	0.133767
W3,W4	0.010347	0.134484

tion, we would like to construct averaged lensing maps $\mathbf{v}^{(QQ)}$, $\mathbf{v}^{(QV)}$, etc. where we only average over differenting assemblies at the same frequency, thereby preserving frequency-dependent information. There are six of these maps (QQ, QV, QW, VV, VW, and WW); the last column of Table I shows the weights used to construct them.

The power spectrum of the longitudinal mode of \mathbf{v} obtained on the cut sky ($Kp0 \cap S10 \setminus ps_2$ cut, which excludes point sources; see Sec. IV A) is shown in Fig. 5.

IV. CROSS-CORRELATION COMPUTATION

A. Sky cuts

In some regions of the sky, particularly the Galactic plane, microwave emission from within the Milky Way and from nearby galaxies dominates over the cosmological signal. For their CMB analysis, the WMAP team removed this signal by (i) masking out a region based on a smoothed contour of the K-band temperature, which they denote “Kp2” [83], and (ii) projecting out of their map microwave emission templates for synchrotron, free-free, and dust emission based on other observations

[84, 85, 86, 87, 88]. Template projection is dangerous for cross-correlation studies involving galaxies because the dust template of Ref. [84] is used to extinction-correct the LRG magnitudes, thus template errors could introduce spurious correlations between the CMB and galaxy maps. Since visual inspection of the uncleared WMAP maps reveals Galactic contamination outside the Kp2-rejected region at all five frequencies, we have used the more conservative Kp0 mask in Sec. VID for our galaxy-temperature correlations. Because the SDSS covers only a small portion of the sky, we speed up the cross-correlation computation by using only WMAP data in the vicinity of the SDSS survey region. We define the “S10” region to consist of those pixels within 10 degrees of the SDSS survey area. The $Kp0 \cap S10$ cut accepts 774,534 HEALPix pixels (10,157 sq. deg.).

When analyzing primary CMB anisotropies, it is customary to mask detected point sources in order to eliminate this spurious contribution to the temperature. For secondary anisotropy studies, the analysis should be done both with and without the point sources because the point sources may correlate with large scale structure, hence naively masking them could lead to misleading results. Therefore for the galaxy-temperature correlation used in Sec. VID we have constructed the $Kp0 \cap S10 \setminus ps$ mask by rejecting all pixels within the WMAP point source mask (with a 0.6 degree exclusion radius around each source). The $Kp0 \cap S10 \setminus ps$ cut accepts 756,078 HEALPix pixels (9915 sq. deg.).

For the lensing analysis, we must use a more conservative mask than Kp0 because the lensing estimator \mathbf{v} is a nonlocal function of the CMB temperature, hence $\mathbf{v}(\hat{\mathbf{n}})$ responds to foreground emission several degrees away from $\hat{\mathbf{n}}$. We have therefore constructed a “Kp05” mask consisting of all pixels within Kp0 that are at least 5 degrees away from the Kp0 boundary; the $Kp05 \cap S10$ mask used for the lensing analysis accepts 753,242 HEALPix pixels (9,878 sq. deg.). We have also constructed a point source-removed version, $Kp05 \cap S10 \setminus ps_2$, in which all pixels within 2 degrees of the point sources are rejected. This mask accepts 598,795 HEALPix pixels (7,853 sq. deg.).

B. Galaxy-convergence correlation

Having constructed the vector field \mathbf{v} , we proceed to compute its cross-correlation C_l^{gv} with the LRG map. We construct the data vector

$$\mathbf{x}^T = (\mathbf{x}_{LRG}^T, \mathbf{x}_{lens}^T) \quad (20)$$

of length $N_{pix}^{(LRG)} + 2N_{pix}^{(CMB)}$, where \mathbf{x}_{LRG} is a vector containing the galaxy overdensities $g = \delta n / \bar{n}$ in each SDSS pixel, and \mathbf{x}_{lens} consists of the two components of \mathbf{v} at each WMAP pixel. (We will suppress the frequency indices (QQ), (QV), etc. on \mathbf{v} for clarity; it is understood that the analysis below is repeated for each pair of

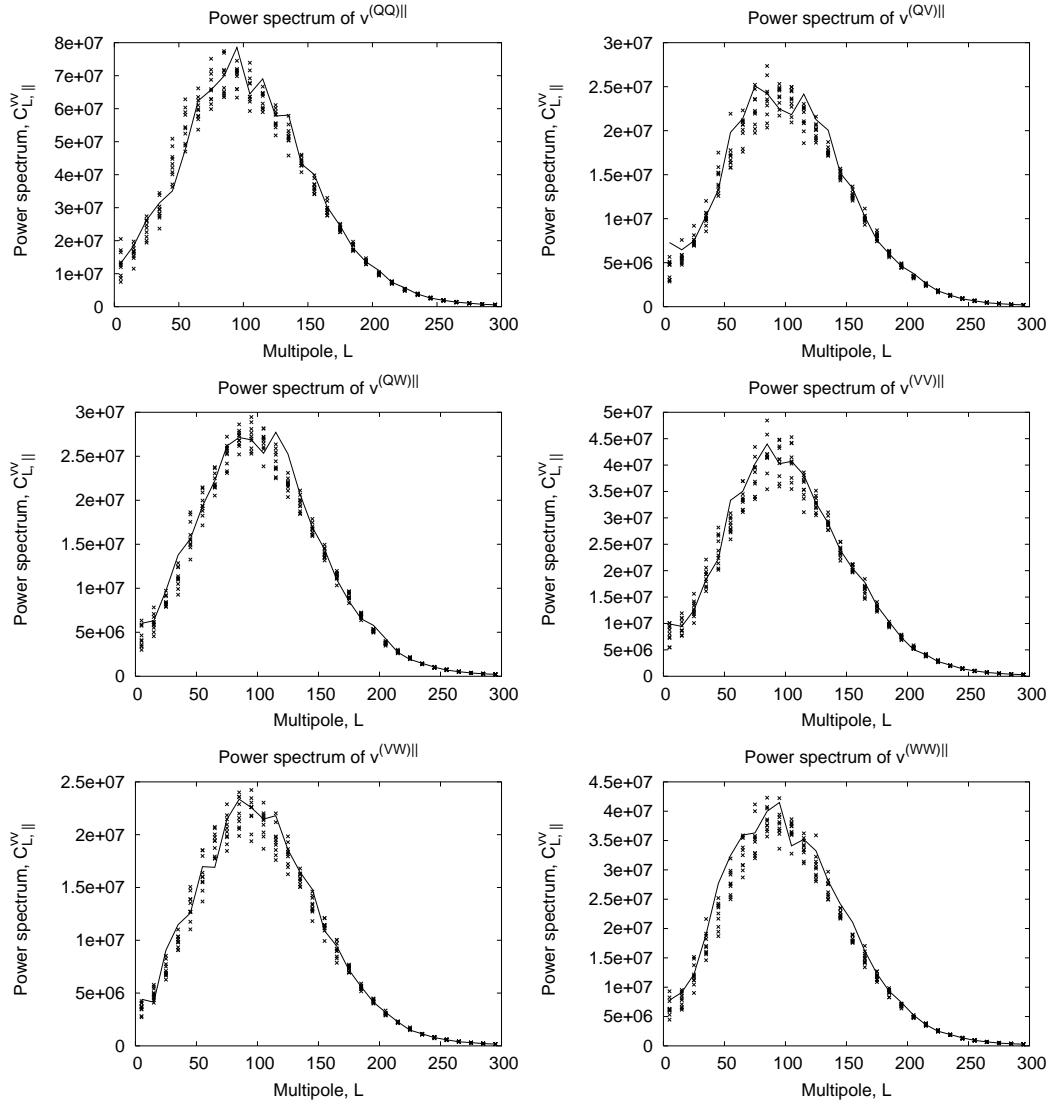


FIG. 5: The longitudinal mode power spectrum $C_{l,\parallel}^{vv}$ within the $Kp05 \cap S10 \setminus ps2$ cut (solid line), for each of the six frequency pairs. The points show 10 simulations (containing no lensing or foregrounds). Since the purpose of this plot is to compare the simulations to the actual data, the sky cut has not been deconvolved, rather we have plotted the average of $|\int \mathbf{v} \cdot \mathbf{Y}_{lm}^{(\parallel)*} d^2\hat{\mathbf{n}}/f_{sky}|^2$ within bands of width $\Delta l = 10$.

frequencies.) The covariance of \mathbf{x} is then:

$$\mathbf{C} \equiv \langle \mathbf{x} \mathbf{x}^T \rangle = \begin{pmatrix} \mathbf{C}^{(LRG)} & \mathbf{C}^{(\times)} \\ \mathbf{C}^{(\times)T} & \mathbf{C}^{(lens)} \end{pmatrix}. \quad (21)$$

The cross-correlation matrix $\mathbf{C}^{(\times)}$ has components:

$$C_{i,jK}^{(\times)} = \sum_{lm} C_l^{gv} Y_{lm}^*(\hat{\mathbf{n}}_i) \mathbf{Y}_{lm}^{\parallel}(\hat{\mathbf{n}}_j) \cdot \hat{\mathbf{e}}_K, \quad (22)$$

where i represents an SDSS pixel index, j is a WMAP pixel index, and $K = \hat{\theta}, \hat{\phi}$ indicates which component of the vector \mathbf{v} is under consideration. We bin the cross-spectrum C_l^{gv} into bands,

$$C_l^{gv} = \sum_A c^A \begin{cases} 1 & l_{\min}(A) \leq l < l_{\max}(A) \\ 0 & \text{otherwise} \end{cases}, \quad (23)$$

and take the c^A as the parameters to be estimated.

In order to construct an optimal estimator for the galaxy-convergence cross-spectrum, we need a prior auto-correlation matrix for the LRGs and for the lensing map. (This is a “prior” in the sense of quadratic estimation theory [89, 90, 91, 92], and has nothing to do with Bayesian priors.) We take a prior of the form

$$C_{ij}^{(LRG)} = \sum_{lm} C_l^{gg} Y_{lm}^*(\hat{\mathbf{n}}_i) Y_{lm}(\hat{\mathbf{n}}_j) + N\delta_{ij}, \quad (24)$$

where C_l^{gg} is the galaxy power spectrum (excluding Poisson noise) and N is the noise variance per pixel. We have taken N to be the reciprocal of the mean number of galaxies per pixel, appropriate for Poisson noise (we use the mean galaxy density per pixel to avoid biases

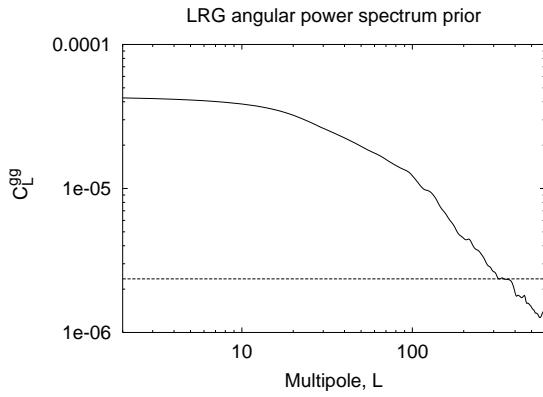


FIG. 6: The prior power spectrum used for the LRG map, obtained by application of a pseudo- C_l estimator to the SDSS scan region. The dashed line shows the Poisson noise.

associated with preferential weighting of pixels containing fewer galaxies). The prior power spectrum $C_l^{(LRG)}$ is determined by application of a pseudo- C_l estimator to the LRG maps; the resulting power spectrum is shown in Fig. 6. We have set $C_l^{gg} = 0.01 \gg C_{l \geq 2}^{gg}$ for $l = 0, 1$ to reject the galaxy monopole (“integral constraint”) and dipole modes from the cross-correlation analysis.

It can be shown [89, 90, 91] that for Gaussian data with small $C^{(\times)}$, the optimal estimator for the c^A would be obtained by taking the unbiased linear combinations of the $\mathbf{x}_{LRG}^T C^{(LRG)-1} \mathbf{P}_A C^{(lens)-1} \mathbf{x}_{lens}$. These estimators are frequently called “QML” or quadratic maximum-likelihood estimators, although they are, strictly speaking, not maximum-likelihood. We do not have the full matrix $C^{(lens)}$, and our only knowledge of this matrix comes from simulations. Thus we have instead constructed, for each lensing map $\mathbf{v}^{\alpha\beta}$, the quadratic combinations:

$$Q_A = \mathbf{x}_{LRG}^T C^{(LRG)-1} \frac{\partial C^{(\times)}}{\partial c^A} \mathbf{x}_{lens}, \quad (25)$$

where A represents a band index. This differs from the QML estimators in that the C^{-1} weighting is applied only to the LRGs, while uniform weighting is applied to the CMB lensing map \mathbf{v} ; thus Eq. (25) can be viewed as a sort of half-QML, half-pseudo- C_l estimator for the cross-spectrum. The expectation value of Q_A can be determined from Eq. (21); it is:

$$\langle Q_A \rangle = \text{Tr} \left[\frac{\partial C^{(\times)T}}{\partial c^B} C^{(LRG)-1} \frac{\partial C^{(\times)}}{\partial c^A} \right] c^B \equiv R_{AB} c^B, \quad (26)$$

which defines the response matrix R_{AB} . Note that unlike the response matrix of the optimal quadratic estimator, R_{AB} is not equal to the Fisher matrix. The trace in Eq. (26) may be computed using a stochastic-trace algorithm:

$$R_{AB} = \langle (C^{(LRG)-1} \mathbf{y})^T \frac{\partial C^{(\times)}}{\partial c^A} \frac{\partial C^{(\times)T}}{\partial c^B} \mathbf{y} \rangle, \quad (27)$$

where \mathbf{y} is a random vector of length $N_{pix}^{(LRG)}$ consisting of ± 1 entries. The vectors in Eqs. (25) and (27) are constructed in pixel space. Harmonic space is used only as an intermediate step in the convolutions required to compute the matrix-vector multiplications, e.g. $\frac{\partial C^{(\times)T}}{\partial c^B} \mathbf{y}$; these are computed by the usual method of converting to harmonic space, multiplying by $\partial C_l^{gg} / \partial c^B$, and converting back to real space. The matrix inverse operations are performed iteratively as described in Appendix B. This method allows us to easily compute estimators for the band cross-powers,

$$\hat{c}^A = [R^{-1}]^{AB} Q_B. \quad (28)$$

While this estimator is manifestly unbiased, we do not know its uncertainty because we do not know the covariance $C^{(lens)}$ of the \mathbf{v} -field. We determine the uncertainty via a Monte Carlo method: we construct random CMB realizations according to the null hypothesis of no lensing in all eight DAs used for the lensing reconstruction, and feed them through the lensing reconstruction pipeline (Sec. III C).

In order to estimate the galaxy bias from the binned cross-power spectrum estimators \hat{c}^A , we need to know the response of each estimator \hat{c}^A to the galaxy bias, $d\langle \hat{c}^A \rangle / db_g$. This is given by

$$\frac{d\langle \hat{c}^A \rangle}{db_g} = [R^{-1}]^{AB} \text{Tr} \left[\frac{dC^{(\times)T}}{db_g} C^{(LRG)-1} \frac{\partial C^{(\times)}}{\partial c^B} \right], \quad (29)$$

which is computed by a stochastic-trace algorithm analogous to Eq. (27). The galaxy- \mathbf{v} correlation matrix $dC^{(\times)T} / db_g$ is

$$\frac{dC_{i,jK}^{(\times)}}{db_g} = \sum_{lm} R_l \frac{dC^{g\kappa}}{db_g} Y_{lm}^*(\hat{\mathbf{n}}_i) \mathbf{Y}_{lm}^{\parallel}(\hat{\mathbf{n}}_j) \cdot \hat{\mathbf{e}}_K, \quad (30)$$

where R_l is the lensing response factor of Eq. (17). If we knew $C^{\mathbf{v}\mathbf{v}}$, it would be optimal to use $C^{\mathbf{v}\mathbf{v}}^{-1}$ weighting, in which case we could simply use $dC^{(\times)T} / db_g$ as a cross-power template with no loss of information. However since we have not calculated $C^{\mathbf{v}\mathbf{v}}$, and our only information on this covariance matrix comes from the ability to generate random realizations of \mathbf{v} , we cannot do this.

C. Simulations

Simulating lensed and unlensed CMB maps is necessary both for verifying the analysis pipeline as well as for determining the optimal weighting of Sec. III D. The general procedure used here is to generate a simulated primary CMB temperature T_{lm} , convergence κ_{lm} , and galaxy density fluctuation δn_{lm} in harmonic space. These are Gaussian random fields and hence it is a straightforward matter to produce random realizations from the power spectra and cross-spectra of T , κ , and δn . After generation of the realization, the primary temperature

and deflection field (generated from the convergence, assuming an irrotational deflection field) are pixelized in HEALPix resolution 10 (12,582,912 pixels of solid angle 11.8 arcmin^2 each). The lensed CMB temperature \tilde{T} is then computed in real space from Eq. (4). Because the “deflected” HEALPix pixels no longer lie on curves of constant latitude, we use a non-isolatitude spherical harmonic transform (see Appendix A; we have used parameters $L' = 6144$ and $K = 11$ since high accuracy is required) to evaluate Eq. (4). The beam convolution relevant to each DA is then applied by converting to harmonic space, multiplying by B_l^α and the pixel window function, and converting back to real space. Finally the simulated CMB temperature field $\tilde{T}(\hat{\mathbf{n}})$ is degraded to HEALPix resolution 9, and appropriate Gaussian “instrument” noise is added independently to each pixel. Note that the resolution 10 pixels are used here to improve the fidelity of the simulation, in particular to ensure that the effects of the elongated HEALPix pixels on the lensing estimator are properly simulated.

A crude model for the WMAP beam ellipticity is incorporated into the simulations as follows. At each point, we have

$$\dot{T}^\alpha(\hat{\mathbf{n}}) = \frac{1}{2} \sum_{\text{side}=A,B} \sum_{lm\mu} B_{l\mu} T_{lm} Y_{lm}^\mu(\hat{\mathbf{n}}) \langle e^{i\mu\psi} \rangle, \quad (31)$$

where the Y_{lm}^μ are spin-weighted spherical harmonics and the beam moments $B_{l\mu}$ are the multipole moments of the beam in instrument-fixed coordinates (with the “North Pole” along the boresight and the $\phi' = 0$ meridian in the scan direction):

$$B_{l\mu} = \sqrt{\frac{4\pi}{2l+1}} \int_{4\pi} B(\hat{\mathbf{n}}_{inst}) Y_{l\mu}^*(\hat{\mathbf{n}}_{inst}) d^2\hat{\mathbf{n}}_{inst}. \quad (32)$$

The average value $\langle e^{i\mu\psi} \rangle$ is taken over the position angles of the instrument when $\hat{\mathbf{n}}$ is scanned. The sum over sides is over the two sides of WMAP. Equation (31) is only an approximation because (i) the two sides of the differencing assembly may not scan each pixel exactly the same number of times, or may have slightly different weights; and (ii) because WMAP is a differential instrument, $\langle \dot{T}^\alpha(\theta, \phi) \rangle$ is also affected by beam-ellipticity effects on other parts of the sky. Since the two beams of a given DA are separated by ~ 140 degrees, this results in an “echo” of a given microwave source at a separation of 140 degrees [40] (and higher-order echoes should also be present); we have neglected these.

We have made several further approximations to Eq. (31) in order to speed up the simulations. First, we have only included the ellipticity modes $\mu = \pm 2$, since these dominate the difference between the azimuthally symmetrized beam and the true beam. The beam ellipticity is thus described by the real and imaginary parts of $B_{l,2}$; recall $B_{l,-2} = B_{l,2}^*$. We have calculated $B_{l,2}$ by taking the non-isolatitude spherical harmonic transform of the WMAP beam maps [64]. Secondly, because the

side A and B beams are approximate mirror-images of each other, we have only considered the component of the beam ellipticity along the scan direction. The component of the beam ellipticity at 45 degrees to the scan direction is suppressed because, to the extent that the side A and B beams are mirror images and scan each pixel the same number of times, this component cancels in Eq. 31 when we sum over the two sides.

Finally, we have used a simple model for the scan pattern $\langle e^{i\mu\psi} \rangle$ for each DA. The WMAP scan pattern is crudely approximated as a rotation around the spacecraft $-Z$ axis, followed by a precession of this axis in a 22.5 degree radius circle around the anti-Sun point, followed by rotation of the anti-Sun point along the ecliptic plane. Relative to the spacecraft $-Z$ axis, the effective number of observations in one rotation is: $N_{obs}(\theta, \phi) = -N_{obs}\langle e^{2i\psi} \rangle = K\delta(\cos\theta - \cos\theta_c)/2\pi$, where θ_c is the angle between the instrument boresight and the spacecraft $-Z$ axis, and K is a constant. We can convert these to harmonic space in the spacecraft coordinates,

$$\begin{aligned} [N_{obs}]_{lm}(-Z) &= K\delta_{m0}Y_{l0}(\theta_c, 0); \\ [N_{obs}\langle e^{2i\psi} \rangle]_{lm}(-Z) &= -K\delta_{m0}Y_{l0}^2(\theta_c, 0). \end{aligned} \quad (33)$$

Averaged over the precession cycle of WMAP, this becomes

$$\begin{aligned} [N_{obs}]_{l0}(\text{av}) &= KP_l(\cos 22.5 \text{ deg})P_l(0)Y_{l0}(\theta_c, 0); \\ [N_{obs}\langle e^{2i\psi} \rangle]_{l0}(\text{av}) &= -KP_l(\cos 22.5 \text{ deg}) \\ &\quad \times P_l(0)Y_{l0}^2(\theta_c, 0). \end{aligned} \quad (34)$$

(The $m \neq 0$ moments vanish.) Once these have been obtained, we may transform back to real space and find $\langle e^{2i\psi} \rangle$ by division. This is then rotated from the ecliptic to the Galactic coordinate system. To speed up computation, the elliptical correction to the beam was only computed on the HEALPix resolution 9 grid whereas the dominant circular part was computed on the resolution 10 grid and then degraded by pixel-averaging.

V. RESULTS

A. Galaxy-convergence cross-spectrum

The individual cross-spectra obtained at different frequencies are shown in Fig. 7. The frequency-averaged cross-spectrum is shown in Fig. 8, both with and without point sources.

B. Amplitude determination

We estimate the bias amplitude b_g by fitting the observed galaxy-convergence cross-spectrum $C_l^{g\kappa}$ to the theoretical model, Eq. (11). We begin by obtaining the covariance matrix $\hat{\Gamma}$ of \hat{c}^A as determined from $M = 50$

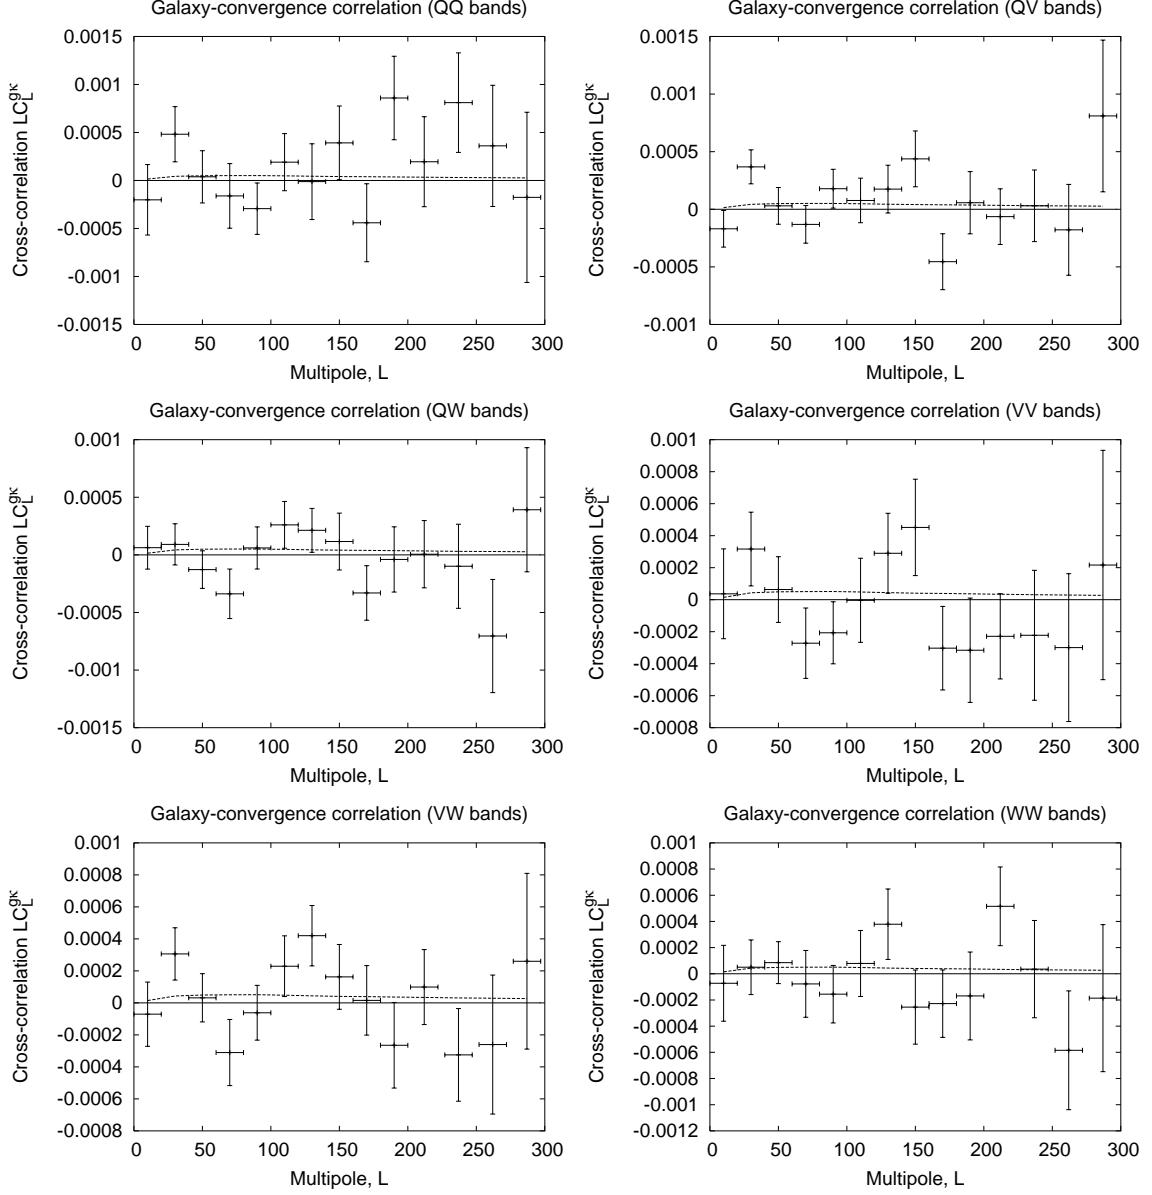


FIG. 7: The galaxy-convergence correlation using the $Kp05 \cap S10 \setminus ps_2$ cut, for each of the six combinations of frequencies. The error bars are strongly correlated across different frequencies. The dashed curve shows the theoretical signal for our best-fit value of $b_g = 1.81$.

simulations:

$$\hat{\Gamma}^{AB} = \frac{1}{M-1} \sum_{i=1}^M (\hat{c}_i^A - \bar{c}^A)(\hat{c}_i^B - \bar{c}^B). \quad (35)$$

(The simulated \hat{c}^A are generated by producing a random realization of the CMB as described in Sec. IV C with no lensing, feeding it through the lensing pipeline, and correlating it against the *real* SDSS LRG map.) The bias is estimated from the weighted average of the observed

cross-powers in our $N = 14$ bands:

$$\hat{b}_g = \frac{\sum_{A=1}^N \hat{c}^A \frac{dc_{th}^A}{db_g} / \hat{\Gamma}^{AA}}{\sum_{A=1}^N \left(\frac{dc_{th}^A}{db_g} \right)^2 / \hat{\Gamma}^{AA}}, \quad (36)$$

where c_{th}^A is the theoretical prediction for the binned cross-spectrum, Eq. (23); this is directly proportional to the bias, $c_{th}^A \propto b_g$. The response dc_{th}^A/db_g is obtained from Eq. (29). It is trivially seen that Eq. (36) is an unbiased estimator of b_g , regardless of the covariance of the \hat{c}^A . Since we are working in harmonic space with bands of width $\Delta l \geq 20 > \Delta\theta^{-1}$, where $\Delta\theta$ is the typical width of the survey region (in radians), the different

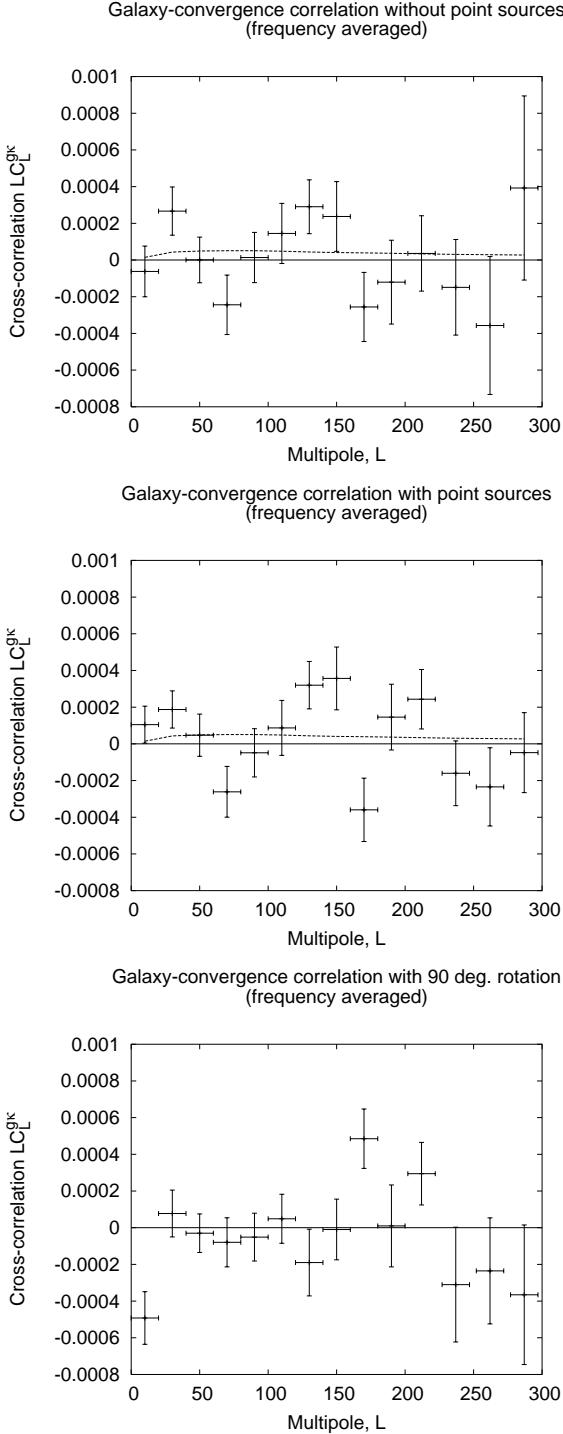


FIG. 8: The galaxy-convergence correlation, using the frequency-averaged $\mathbf{v}^{(TT)}$ map. The top panel shows the correlation using the $\text{Kp05} \cap \text{S10} \setminus \text{ps2}$ mask (which rejects point sources). The middle panel shows the correlation using the $\text{Kp05} \cap \text{S10}$ mask, which does not reject point sources. The bottom panel shows the correlation with the $\mathbf{v}^{(TT)}$ field rotated by 90 degrees; this should be zero in the absence of systematics (see Sec. VI A). The error bars are from simulations as described in Sec. V B. The dashed curve shows the theoretical signal for $b_g = 1.81$.

l -bands are very weakly correlated, so we have not attempted to further optimize the relative weights of the various cross-power estimators \hat{c}^A .

The most obvious way to estimate the uncertainty in \hat{b}_g by noting that Eq. (36) is a linear function of the c^A , and substituting in the covariance matrix $\hat{\Gamma}^{AB}$ of the $\{c^A\}$:

$$\sigma(\hat{b}_g; \text{incorrect}) = \frac{\sum_{A,B=1}^N \hat{\Gamma}^{AB} \frac{dc_{th}^A}{db_g} \frac{dc_{th}^B}{db_g} / \hat{\Gamma}^{AA} \hat{\Gamma}^{BB}}{\left[\sum_{A=1}^N \left(\frac{dc_{th}^A}{db_g} \right)^2 / \hat{\Gamma}^{AA} \right]^2}. \quad (37)$$

This calculation is incorrect for finite number $M < \infty$ of simulations because it neglects the fact that the $\hat{\Gamma}^{AB}$ are themselves random variables. One approach to the problem is to take a sufficiently large number of simulations M that the error in Eq. (37) becomes negligible. The difficulties in this approach are that it could be very computationally intensive; we do not know whether M simulations are “enough” unless we try even larger values of M to check convergence. An alternative method, which we have used here, is to run an additional $M' = 50$ simulated realizations of $\{\hat{c}^A\}$ (identical to those used to compute $\hat{\Gamma}^{AB}$ except for the random number generator seed), compute the bias \hat{b}_g from them, and then compute their sample variance. The resulting error bars can be analyzed using the well-known Student’s t -distribution. The “ 1σ ” t error bars (which have 49 degrees of freedom) obtained by this method are shown in Table II. The mean bias values obtained from these 50 random realizations are shown in the “random” column in the table. Also shown in Table II (in the “foreground” columns) are the results obtained by feeding the Galactic foreground templates of Sec. VI E 2 through the lensing pipeline and correlating these with the real LRG map.

In the case of the $\text{Kp05} \cap \text{S10}$ cut (last column in Table II), which does not reject point sources, the $\mathbf{v}^{(QQ)}$, $\mathbf{v}^{(QV)}$, and $\mathbf{v}^{(QW)}$ maps have power spectra that are boosted significantly by point source contamination (see Sec. VID 1). Therefore, even if the correlation of the point sources with the galaxies can be neglected, the error $\sigma(b_g)$ obtained in these bands for the $\text{Kp05} \cap \text{S10}$ mask is probably underestimated, as noted in the table.

The $\hat{\chi}^2$ values for fits to zero signal are shown in Table III. These are obtained using the 14 band cross-power spectra (Fig. 8), and the 14×14 covariance matrix is obtained from 100 simulations,

$$\hat{\chi}^2 = [\hat{\Gamma}^{-1}(100 \text{ sims})]_{ABC} c^A c^B. \quad (38)$$

Because the number of simulations is finite, there remains some noise in this covariance matrix and this must be taken into account in interpreting the $\hat{\chi}^2$. In particular, the $\hat{\chi}^2$ variable does not exactly follow the standard χ^2 (so we have denoted it with a hat). The distribution and p -values can, however be calculated as described in Ref. [93], Appendix D. As noted previously, the errors for the $\text{Kp05} \cap \text{S10} \setminus \text{ps2}$ mask in the QQ, QV, and QW combinations are suspect.

TABLE II: The bias b_g estimated using several frequency ranges (first column). The second column gives the bias b_g obtained with the point sources removed ($Kp05 \cap S10 \setminus ps_2$ cut); these are the numbers that should be thought of as our “result.” The uncertainties are 1σ with t errors (49 d.o.f.); the first error given used the elliptical beam simulations, the error in parentheses is obtained from circular beam simulations (Sec. VI C). The “TT” frequency combination is a weighted average of QV, QW, VV, VW, and WW. The column labeled “foreground” show the bias b_g obtained by correlating the \mathbf{v} maps derived from the Galactic foreground maps of Sec. VI E 2 against the LRG map. Similarly, the column labeled “random” show the bias b_g obtained using the average of the 50 random realizations of \mathbf{v} in place of the *WMAP*-derived \mathbf{v} map. The column labeled “TT weight” shows the bias determined by using the same weights as a function of l as for the TT frequency combination in Eq. (36); of course this has no effect on $b_g(TT)$. The final column, labeled “ $Kp05 \cap S10$ ”, is the bias obtained without the point source cut.

Frequency	Bias, b_g	foreground	random	TT weight	b_g , $Kp05 \cap S10$
QQ	$+3.62 \pm 4.48 (\pm 4.33)$	-0.001640	-0.34 ± 0.63	$+3.13 \pm 4.55$	$+6.30 \pm 3.93^a$
QV	$+3.10 \pm 2.19 (\pm 2.00)$	-0.000001	$+0.07 \pm 0.31$	$+3.39 \pm 2.29$	$+3.93 \pm 1.99^a$
QW	$-0.11 \pm 2.53 (\pm 2.35)$	-0.000480	$+0.06 \pm 0.36$	-0.20 ± 2.56	$+0.48 \pm 2.43^a$
VV	$-1.41 \pm 2.95 (\pm 3.13)$	-0.000737	$+0.88 \pm 0.42$	$+0.34 \pm 3.04$	$+0.58 \pm 2.35$
VW	$+2.63 \pm 2.56 (\pm 2.11)$	-0.000617	$+0.35 \pm 0.36$	$+2.58 \pm 2.62$	$+2.51 \pm 2.03$
WW	$+0.23 \pm 3.11 (\pm 2.71)$	-0.000817	$+0.20 \pm 0.44$	$+0.65 \pm 3.22$	-0.75 ± 2.63
TT	$+1.81 \pm 1.92 (\pm 1.72)$	-0.000449	$+0.23 \pm 0.27$	$+1.81 \pm 1.92$	$+2.43 \pm 1.58^a$

^aBecause of point sources, the \mathbf{v} maps from these bands contain excess power in the $Kp05 \cap S10$ region. Thus the simulation error bars shown here are likely underestimates.

TABLE III: The $\hat{\chi}^2$ values obtained for fits to zero signal from the galaxy-convergence cross-spectrum. The first column shows the results with the $Kp05 \cap S10$ mask; the second using the $Kp05 \cap S10 \setminus ps_2$ mask; and the third using the $Kp05 \cap S10 \setminus ps_2$ mask with 90 degree rotation of \mathbf{v} . The χ^2 has 14 degrees of freedom (the 14 l -bins shown in Fig. 8) and the covariance matrices were obtained from the 100 simulations described in Sec. V B. As described in the text, the finite number of simulations means that the expectation value of the $\hat{\chi}^2$ is not 14 but is larger due to uncertainty in the covariance matrix; the mean of the $\hat{\chi}^2$ distribution is 16.5 and the standard deviation is 6.8. We have also given the cumulative probability distributions.

Freq.	$Kp05 \cap S10$		$Kp05 \cap S10 \setminus ps_2$		$Kp05 \cap S10 \setminus ps_2$ +90 degrees	
	$\hat{\chi}^2$	$P(< \hat{\chi}^2)$	$\hat{\chi}^2$	$P(< \hat{\chi}^2)$	$\hat{\chi}^2$	$P(< \hat{\chi}^2)$
QQ	51.51 ^a	0.9996 ^a	16.24	0.55	28.24	0.940
QV	37.01 ^a	0.990 ^a	27.51	0.931	21.89	0.81
QW	16.74 ^a	0.58 ^a	11.74	0.26	21.58	0.80
VV	19.08	0.70	11.86	0.27	24.85	0.89
VW	25.30	0.90	20.28	0.75	20.06	0.74
WW	11.54	0.25	11.36	0.24	21.82	0.81
TT	26.90 ^a	0.922 ^a	19.97	0.74	30.71	0.963

^aBecause of point sources, the \mathbf{v} maps from these bands contain excess power in the $Kp05 \cap S10$ region. Thus the simulation error covariance matrices are likely underestimates, and hence the $\hat{\chi}^2$ and $P(< \hat{\chi}^2)$ values are suspect.

VI. SYSTEMATIC ERRORS

A. Ninety-degree rotation test

One of the standard systematics tests in weak lensing studies using galaxies as sources has been to rotate all of the galaxies by 45 degrees and look for a shear signal. The 45 degree rotation is used because it interconverts

E and B modes, and in the absence of systematics there should be no B -mode signal. In the case of CMB lensing using the vector estimator \mathbf{v} , the analogous test is to rotate \mathbf{v} by 90 degrees (thereby interchanging the longitudinal and transverse parity modes). This rotated map can be fed through the cross-correlation pipeline in place of the original \mathbf{v} . In the absence of systematics, this gives zero signal; the error bars need not be the same as for the longitudinal modes, but they can still be determined from simulations as described in Sec. V B. The cross-spectrum is shown in Fig. 8 and the $\hat{\chi}^2$ values in Table III.

The lowest- l point in the rotated cross-spectrum (Fig. 8) is 3.4σ negative. It is difficult to assess the significance of this anomaly since it is an *a posteriori* detection ($p = 0.00127$ for the two-tailed t -distribution); in any case, it is responsible for the relatively high $\hat{\chi}^2$ value ($p = 0.037$) in the $Kp05 \cap S10 \setminus ps_2 + 90$ degree column of Table III. It is unlikely that this correlation $\langle g^* v_{lm}^{(TT)\perp} \rangle$ represents any real astrophysical or cosmological effect, since it violates parity. This anomaly is also distinct from the much-discussed “low quadrupole” observed by *WMAP*, since the former is based on a high-pass filtered CMB map with power coming predominately from CMB modes with $l \sim \text{few} \times 10^2$. Another possible explanation would be some source of excess power in the \mathbf{v} map at low l , which would increase the error bar relative to simulations and thus lower the statistical significance of this point. However if we take the \mathbf{v} maps, and compute the un-deconvolved power spectrum

$$P = \sum_{l < 20} \sum_{m=-l}^l \left| \sum_{\hat{\mathbf{n}} \in Kp05 \cap S10 \setminus ps_2} \mathbf{v}(\hat{\mathbf{n}}) \cdot Y_{lm}^{\perp*}(\hat{\mathbf{n}}) \right|^2 \quad (39)$$

for both the real \mathbf{v} map and the 100 simulated maps, we find that the real map has the 28th highest value of P out of 101 maps, i.e. there is no evidence for excess power. If

this point is due to some systematic, it must be present at all three frequencies, since this point is negative by at least 1σ in all of the frequency combinations except QQ, where the binned $C_l^{g\kappa}$ from the rotated map at $l < 20$ is $(1.0 \pm 3.5) \times 10^{-5}$.

It is thus difficult to explain the lowest- l point in Fig. 8 point in terms of any systematic error. The true test for whether this is in fact just a statistical fluctuation is to wait for the error bars to become smaller with future *WMAP* data and see whether this point becomes more significant or goes away, and in the former case whether it exhibits a frequency dependence.

B. End-to-end simulation

Another important systematic test is to verify, in an “end-to-end” simulation, that the lensing estimator and $C_l^{g\kappa}$ cross-spectrum estimator are calibrated properly. This can be done as follows. We run 50 simulations in which simulated Gaussian g and κ maps are generated with the cross-spectrum appropriate for $b_g = 1$. The κ map has the power spectrum $C_l^{\kappa\kappa}$ expected for a Λ CDM cosmology, while the g map is constructed from $g_{lm} = C_l^{g\kappa} \kappa_{lm} / C_l^{\kappa\kappa}$. In principle one could add additional noise to g to boost its power spectrum to match the observed C_l^{gg} , but there is no reason to do this as it increases the number of simulations required and has no effect on the calibration. The κ maps are then used to generate lensed CMB maps using the simulation code described in Sec. IV C, and the output temperature maps $\hat{T}_{Q1} \dots \hat{T}_{W4}$ fed through the lensing reconstruction pipeline and then the $C_l^{g\kappa}$ estimator. Finally, we estimate the bias in each simulation using Eq. (36). This output b_g is the calibration factor $1 + \zeta$ appropriate for cross-correlation studies.

The calibration factors obtained from this procedure are shown in Table IV; the table reveals that the lensing pipeline is calibrated at the $\sim 20\%$ level. Calibration factors of this order have been observed in previous simulations [76, 94] and have been investigated analytically [76, 77, 95, 96], where the main effect has been the non-linear lensing effects (i.e. the order κ^2 and higher terms that have been dropped in the Taylor expansion, Eq. 5). In our case, the calibration error may also have a contribution from the elliptical beam. In any case, the calibration errors are not significant at the level of the current data (i.e. no detection).

C. Beam effects

We have used a crude model for the *WMAP* beam ellipticity. An incorrect model for the beam can have three effects on the galaxy-convergence cross spectrum and hence on the bias determination: it can (i) produce a shear calibration bias in the \mathbf{v} map (which may depend on the wavenumber l and orientation of the convergence

TABLE IV: The calibration factors obtained via end-to-end simulations, for the Kp05 \cap S10\ps2 mask and various combinations of frequencies. Error bars are 1σ , t -distributed with 49 degrees of freedom.

Frequencies	Calibration factor, $1 + \zeta$
QQ	1.17 ± 0.13
QV	1.24 ± 0.09
QW	1.19 ± 0.07
VV	1.31 ± 0.11
VW	1.15 ± 0.07
WW	0.92 ± 0.10
TT	1.18 ± 0.07

mode in question, and may vary across the sky as the effective beam varies); (ii) modify the noise covariance matrix of \mathbf{v} ; and (iii) introduce artifacts (i.e. biases) in the \mathbf{v} map because it invalidates the assumption that the signal is statistically isotropic. The calibration problem is obviously of concern for attempts to do precision cosmology with lensing. However since we do not have a detection, the only effect of the calibration bias is to affect our upper limits on the lensing signal. The change in the noise covariance of the lensing map \mathbf{v} is potentially more serious because it can alter the variance of our cross-spectrum estimator and hence affect the statistical significance of any lensing detection. Because artifacts in the \mathbf{v} map do not correlate with the galaxy distribution, they are essentially also a source of spurious power, and for cross-correlation measurements they are only a concern if their power spectrum is comparable to that of the noise.

The effect of the beam on the noise covariance can be addressed as follows. We have re-computed the uncertainties $\sigma(b_g)$ (see Sec. V B) using 50 simulations with a circular beam instead of our elliptical beam model. The uncertainties are shown in parentheses in Table II. They are at most 20% different from the error estimates obtained from the elliptical beams (but this may not be significant because the $\sigma(b_g)$ values from simulations are themselves drawn from a random distribution – namely, the square root of a χ^2 distribution with 49 degrees of freedom – and hence have an uncertainty of $1/\sqrt{2 \cdot 49} \approx 10\%$). Since replacing our model of the beam ellipticity with the inferior model of a circular beam has only a $< 20\%$ effect on $\sigma(b_g)$, it is doubtful that $\sigma(b_g)$ would be altered by more than this by use of an improved beam model.

Finally, we come to the issue of the calibration. Our estimator for the bias was constructed assuming a circular beam, and given that the true beam is not circular, we expect that it may be mis-calibrated, i.e. $\langle \hat{b}_g \rangle = (1 + \zeta) \hat{b}_g$, where ζ is the calibration bias. At present, the best way to test for such a bias is via simulations, such as those of Sec. VI B. There we found a calibration bias of $\zeta = 0.18 \pm 0.07$, which is not important at the level of the present data.

D. Extragalactic foregrounds

The lensing estimator of Eq. (14) will respond not only to real lensing signals, but to any other perturbations of the CMB. Of greatest concern is the contamination from extragalactic foregrounds, which may induce spurious correlation of the lensing estimator with the galaxy distribution since the extragalactic foregrounds (tSZ and point sources) are expected to correlate with large scale structure. The presence of the extragalactic foregrounds causes the observed temperature $\hat{T}(\hat{\mathbf{n}})$ to be incremented by some amount $\Delta T^\alpha(\hat{\mathbf{n}})$. Assuming that the extragalactic foregrounds are not correlated with the primary CMB or with instrument noise, this causes the expectation value of \mathbf{v} (averaged over primary CMB and noise realizations) to be incremented by (compare to Eq. 15)

$$\Delta\langle v_{lm}^{\alpha\beta(\parallel)} \rangle = (-1)^m \sum_{l'l''} \mathcal{K}_{ll'l''} \sum_{m'm''} \binom{l \ l' \ l''}{-m \ m' \ m''} \times \frac{\Delta T_{l'm'}^\alpha \Delta T_{l''m''}^\beta + \Delta T_{l'm'}^\beta \Delta T_{l''m''}^\alpha}{2}. \quad (40)$$

Thus the possible source of contamination of the convergence-galaxy correlation signal $C_l^{g\mathbf{v}}$ is the correlation of the LRGs with quadratic combinations of the foreground temperature. In the case of the tSZ foreground, the contribution to Eq. (40) can be broken up into a “single-halo” term in which the two factors of ΔT come from the same halo, and a “two-halo” term in which the two factors of ΔT come from different halos. The single-halo term exists even if the tSZ halos are Poisson-distributed, whereas the two-halo term acquires a nonzero value only from clustering of the halos. Much of this section will be devoted to an investigation of the properties of the quadratic combinations in Eq. (40) and an assessment of their magnitude. Unfortunately, we will see that this does not result in useful constraints on the foreground contamination to our measurement of b_g .

1. Point sources

It is readily seen that point sources are a major contribution to the power spectrum of \mathbf{v} , especially in the lower-frequency bands. This can be seen from Fig. 9, in which the power spectrum $C_{l,\parallel}^{vv}$ is shown in the Kp05 \cap S10 region (in which point sources are not masked). The $\mathbf{v}^{(QQ)}$, $\mathbf{v}^{(QV)}$, and $\mathbf{v}^{(QW)}$ maps are heavily contaminated while for the higher-frequency maps point sources are subdominant to CMB fluctuations. While the contribution to the \mathbf{v} autopower in Kp05 \cap S10 is large, we are interested here in whether – and at what frequencies – the point source contribution to $\mathbf{v}^{\alpha\beta}$ correlates with the LRG map when the Kp05 \cap S10\ps2 mask (which masks point sources) is used.

A single point source with frequency-dependent flux $F_a(\nu)$ (in units of blackbody μK sr) at position $\hat{\mathbf{n}}_a$ will

produce a spurious contribution to the temperature of

$$\Delta T_{lm}^\alpha = F_a(\nu^\alpha) Y_{lm}^*(\hat{\mathbf{n}}_a). \quad (41)$$

Plugging this into Eq. (40), we find that the shift in $\langle \mathbf{v} \rangle$ is:

$$\Delta\langle v_{lm}^{\alpha\beta(\parallel)} \rangle = F_a(\alpha) F_a(\beta) Y_{lm}^*(\hat{\mathbf{n}}_a) r_{PS}(l), \quad (42)$$

where the response function $r_{PS}(l)$ is:

$$r_{PS}(l) = \sqrt{\frac{(2l'+1)(2l''+1)}{4\pi(2l+1)}} \binom{l \ l' \ l''}{0 \ 0 \ 0} \mathcal{K}_{ll'l''}. \quad (43)$$

In Table V, we show the product $F_a(\nu_1)F_a(\nu_2)$ for several possible point source spectra. Note that for the steep spectra characteristic of WMAP point sources ($\alpha \sim 0.0$), the contamination of $\mathbf{v}^{(QQ)}$ and $\mathbf{v}^{(QV)}$ is far greater than contamination of the higher-frequency lensing maps. Therefore these lower-frequency bands are a useful test of point source contamination of the galaxy-convergence correlation. The dependence of the estimated b_g on the combination of frequencies, $F_a(\nu_1)F_a(\nu_2)$, is shown in Fig. 10. If there were point source contamination of our measurement with spectral index $\alpha = 0.0$, the points in Fig. 10 would be expected to fall roughly along a line $b_g \propto F_a(\nu_1)F_a(\nu_2)$; this is only rough because the weighting of different l bins is slightly different at different frequencies in Eq. (36), and the contaminating signal need not have the same angular dependence as the galaxy-convergence correlation. If we re-calculate the six frequency combinations $b_g(QQ)\dots b_g(WW)$ using the same weighting of different l bins as for the $b_g(TT)$ measurement, we get the values in the column in Table II labeled “TT weight.” Assuming a synchrotron-like spectrum for the point sources, a correlated least-squares fit of the form

$$b_g(\nu_1\nu_2) = b_g^{(0)} + \Delta b_g^{(PS)}(TT) \frac{F_a(\nu_1)F_a(\nu_2)}{\langle F_a F_a \rangle(TT)} \quad (44)$$

to the various frequency combinations will return for $b_g^{(0)}$ a point-source-marginalized measurement of the bias, and for $\Delta b_g^{(PS)}(TT)$ a measurement of the point-source contamination to the unmarginalized $b_g(TT)$. The results of such a fit are $b_g^{(0)} = 0.58 \pm 2.36$ and $\Delta b_g^{(PS)}(TT) = 0.73 \pm 1.18$; the two measurements are of course anti-correlated with correlation coefficient $\rho = -0.63$. There is thus no evidence for point source contamination, although the statistical errors are too large to definitively say whether such contamination is present at the level of the signal. It would be useful to have lower-frequency information here in order to improve the constraints, however this is not possible as the CMB multipoles used in our analysis ($l \sim 350$) are not resolved by WMAP K- and Ka-band differencing assemblies.

The contamination in $C_l^{g\mathbf{v}}$ from a point source can also be estimated from angular information since the point

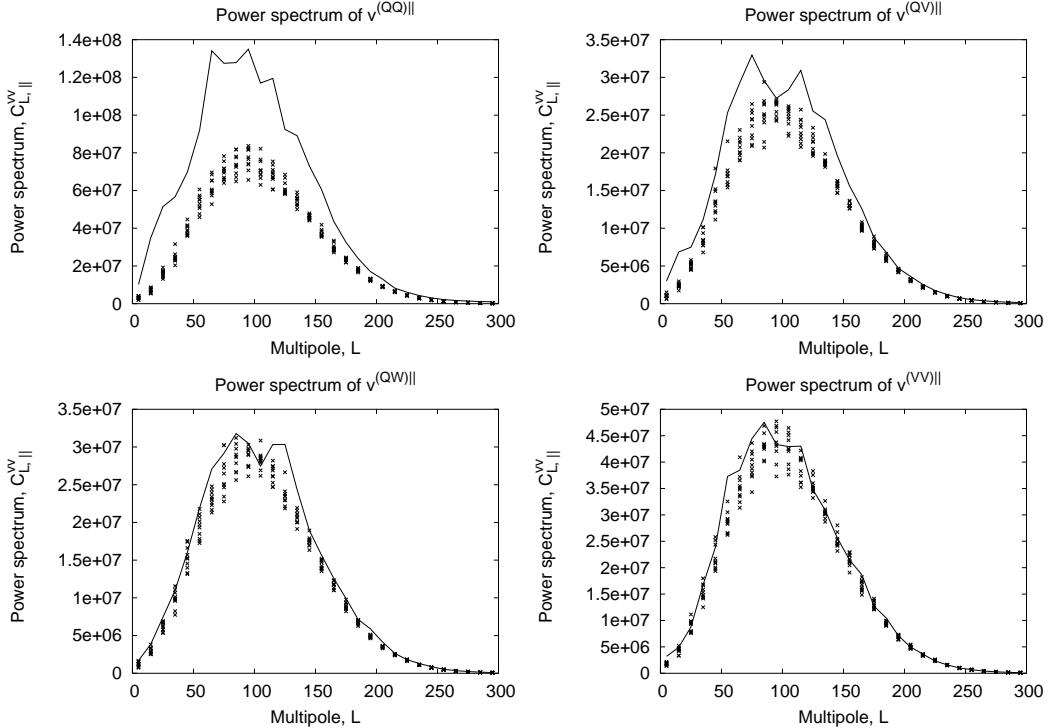


FIG. 9: The cut-sky power spectrum of the longitudinal mode of \mathbf{v} , $C_{l,\parallel}^v$; the power spectrum of the actual map is shown as the solid line, while the points are from 10 simulations. This figure is identical to Fig. 5, except that we have used the $Kp0 \cap S10$ cut (i.e. in this figure the point sources are not masked), and we have not shown the VW or WW spectra (which do not differ significantly from $Kp0 \cap S10 \setminus ps$). Because of contamination by point sources, the power in the QQ and QV maps is greater than in the simulations or the point source-cut maps.

sources have an angular dependence (roughly Poisson) unlike that of the CMB. Equations (41) and (42) give

$$\Delta C_l^{g\kappa(TT)} = \frac{r_{PS}(l)}{R_l} \left[\sum_{\alpha\beta} a_{\alpha\beta}^{(TT)} \frac{F(\alpha)F(\beta)}{F(Q)^2} \right] \times F(Q) \Delta C_l^{gT(Q)}, \quad (45)$$

where $\Delta C_l^{gT(Q)}$ is the point source-induced galaxy-temperature cross spectrum in Q-band. If take a typical spectrum $\alpha = 0.0$ and a flux $F(Q) = 1$ Jy (i.e. the brightest point sources not excluded by the WMAP point source mask [83, 97]), we find

$$\left[\sum_{\alpha\beta} a_{\alpha\beta}^{(TT)} \frac{F(\alpha)F(\beta)}{F(Q)^2} \right] F(Q) = 6.7 \times 10^{-3} \mu\text{K sr}. \quad (46)$$

The ratio $r_{PS}(l)/R_l$ is plotted in Fig. 11. Of course, not all of the point sources have $F(Q) = 1$ Jy, but this is the worst-case scenario since $\Delta C_l^{g\kappa} \propto F \Delta C_l^{gT}$, hence if the galaxy-temperature cross-spectrum is coming from fainter sources the contamination will be even less. (This scaling with F occurs because the spurious contribution to the galaxy-convergence cross spectrum is quadratic in the flux whereas the contribution to the galaxy-temperature cross spectrum is linear.) We have com-

puted the Q-band galaxy-temperature cross-spectrum using a QML estimator [98] on the $Kp0 \cap S10$ cut; if we take this cross-spectrum, and assume that at $l > 60$ the cross-spectrum is entirely due to point sources with the synchrotron spectrum, the derived contamination to the galaxy-convergence spectrum is as shown in Fig. 12(a). One can also use the difference between Q and W-band cross-spectra; in this case, if a synchrotron spectrum with $\alpha = 0.0$ for the point sources is assumed, the difference $C_l^{gT}(Q) - C_l^{gT}(W)$ must be multiplied by 1.295 to recover the $C_l^{gT}(Q)$ used in Eq. (45). This result is shown in Fig. 12(b); the error bars (obtained from simulations) are now smaller at low l because the CMB fluctuations are suppressed. The point-source induced error in the bias can be computed from the data in Fig. 12(b) by plugging these $\Delta C_l^{g\kappa}$ values into Eq. (36); the result is $\Delta b_g^{(PS)} = -0.14 \pm 0.51$.

We conclude that the point-source contamination to $C_l^{g\kappa}$ is at most of the same order as the signal in this range of multipoles. If one ignores correlations between distinct point sources so that Eq. (45) is valid and assumes the $\alpha = 0.0$ spectrum, then Fig. 12(b) suggests that the point source contamination is less than the observed signal.

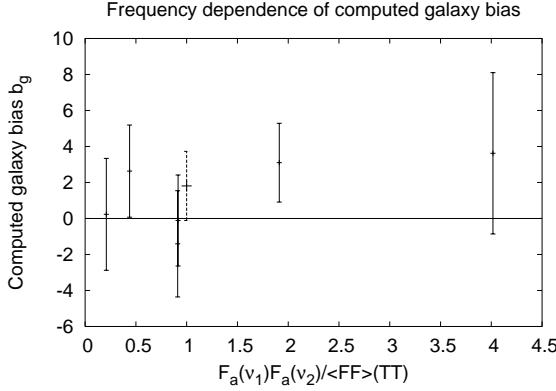


FIG. 10: The dependence of the computed galaxy bias b_g on frequency. The horizontal axis is $F_a(\nu_1)F_a(\nu_2)/0.249$, computed for a typical synchrotron spectrum ($\alpha = 0.0$); the normalization is chosen so that the frequency-averaged TT result has $F_a(\nu_1)F_a(\nu_2)/0.249 = 1$. The horizontal line is zero, and the point with dashed error bars is TT . (This point is of course the average of the other data points and contains no additional information.) The error bars are correlated; the χ^2 for a frequency-independent bias is 6.31 with 5 degrees of freedom.

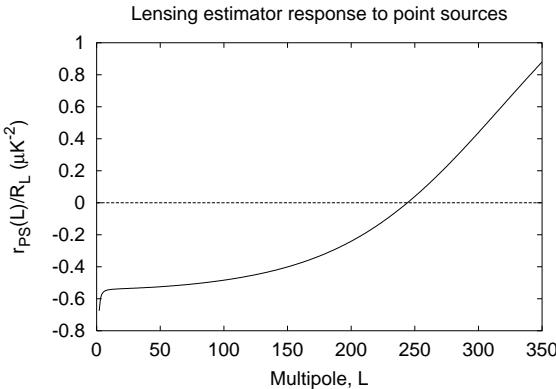


FIG. 11: The ratio r_{PS}/R_l describing the response of the lensing estimator \mathbf{v} to contamination from Poisson-distributed point sources. Note that on large scales this is negative, i.e. regions with more point sources are interpreted by the lensing estimator as regions of negative convergence (underdense regions). At high l the situation is reversed.

2. Thermal Sunyaev-Zel'dovich effect

In the case of the tSZ effect, the frequency dependence is exactly known; see Table V. Unfortunately this frequency dependence is extremely weak in the WMAP bands, with $F_a(\nu_1)F_a(\nu_2)$ varying by only a factor of 0.667 from the QQ to WW bands. Therefore we must resort to the angular dependence to separate tSZ from lensing.

In order to perform an analysis similar to that of Sec. VID 1, we need to determine or place a bound on the galaxy-SZ power cross-spectrum, and determine the

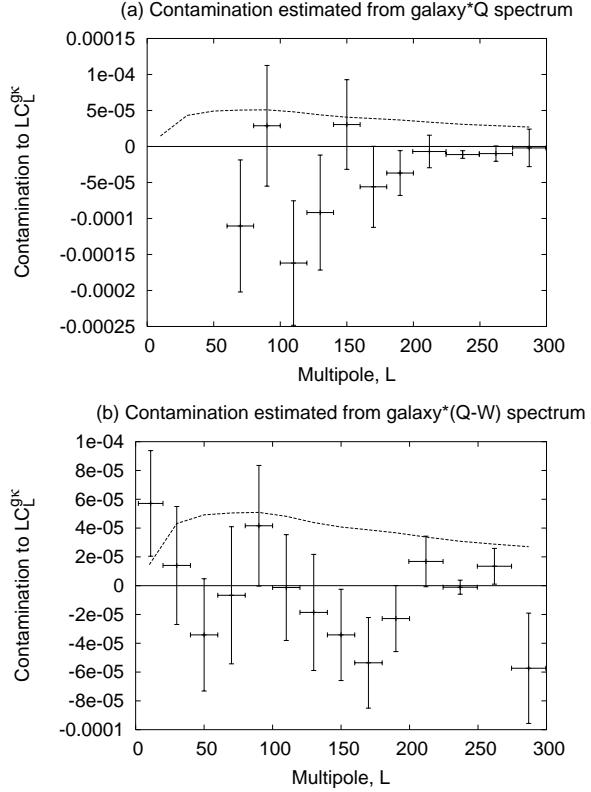


FIG. 12: (a) The contamination to the galaxy-convergence cross-spectrum due to point sources, based on the Q-band galaxy-temperature cross-spectrum in the $Kp0 \cap S10$ region (see text). The dashed line is the best-fit signal from Sec. V B, with $b_g = 1.81$. We have removed the $l < 60$ points since these may contain ISW signal. (b) The same figure, except derived from the difference of galaxy-temperature cross spectra in Q and W bands. This cancels any contribution from ISW and reduces cosmic-variance errors at low l .

maximum flux $|F|$ of a tSZ source (we use $|F|$ since tSZ sources have negative net flux in the WMAP bands). In principle, tSZ haloes can be resolved, however since our lensing estimator uses information from multipoles $l \sim 350$ (physical wavenumber $k/a = 0.7 \text{ Mpc}^{-1}$ at $D_{AP} \geq 0.5 \text{ Gpc}$), the haloes will not be resolved. This argument of course only applies to the single-halo contribution to Eq. (40). The flux from a tSZ source is

$$\begin{aligned} F &= -f_b T_{CMB} \frac{\sigma_T}{D_{AP}^2} \frac{k_B \langle T_e \rangle}{m_e c^2} \frac{M}{\mu_e m_p} \left(4 - x \coth \frac{x}{2} \right) \\ &= -8 \times 10^{-4} \mu\text{K sr} \left(\frac{M \langle T_e \rangle}{10^{16} M_\odot \text{ keV}} \right) \left(\frac{f_b / \mu_e}{0.18} \right) \\ &\quad \times \left(\frac{D_{AP}}{\text{Gpc}} \right)^{-2} \left(4 - x \coth \frac{x}{2} \right), \end{aligned} \quad (47)$$

where σ_T is the Thomson cross section, k_B is Boltzmann's constant, $f_b = \Omega_b / \Omega_m$, m_e and m_p are the electron and proton masses, $\langle T_e \rangle$ is the mean electron temperature, $\mu_e m_p$ is the baryonic mass per free electron, D_{AP} is the physical angular diameter distance, M is the total mass

TABLE V: The point source frequency-dependent response $F_a(\nu_1)F_a(\nu_2)$ for several point source spectra. Included are three power-law spectra, $T_{\text{antenna}} \propto \nu^{\alpha-2}$ ($\alpha = -1.0, 0.0$, and $+0.5$) and a nonrelativistic tSZ spectrum. We have normalized to unit response in $\mathbf{v}^{(QQ)}$. The “TT” frequency combination is an average over the other pairs of frequencies weighted by $a_\alpha^{(TT)}\beta$.

Bands	$F_a(\nu_1)F_a(\nu_2)$			
	$\alpha = -1.0$	$\alpha = 0.0$	$\alpha = +0.5$	tSZ
QQ	1.000	1.000	1.000	1.000
QV	0.320	0.476	0.581	0.946
QW	0.099	0.228	0.345	0.817
VV	0.102	0.227	0.338	0.895
VW	0.032	0.109	0.200	0.773
WW	0.010	0.052	0.119	0.667
TT	0.137	0.249	0.350	0.838

of the halo, and $x = h\nu/k_B T_{CMB} = \nu/57$ GHz. The frequency-dependent factor $4 - x \coth(x/2)$ ranges from 1.91 (Q band) to 1.57 (W band). For tSZ sources that are physically associated with the LRGs (distance $z \geq 0.2$ or $D_{AP} \geq 0.5$ Gpc) we will have $|F| < 6 \times 10^{-3} \mu\text{K sr}$ even for extremely massive ($M\langle T_e \rangle = 10^{16} M_\odot$ keV) haloes. If we estimate contamination to the galaxy-convergence correlation in analogy to Eq. (45), we find

$$\Delta C_l^{g\kappa}(\text{tSZ, 1 halo}) = \frac{r_{PS}(l)}{R_l} F \Delta C_l^{gT}; \quad (48)$$

if we take $|F| = 6 \times 10^{-3} \mu\text{K sr}$, the limits on the contamination from the Q-band correlation C_l^{gT} are similar to the limits for point sources from $C_l^{gT}(Q)$ (cf. Fig. 12a). Unfortunately, like our similar analysis for point sources, this analysis of the tSZ contamination does not tell us anything new, since if the contamination from tSZ in our b_g measurement were large compared to the statistical error of ± 1.92 , we would have measured the wrong b_g . Worse, it only applies to the single-halo contribution in Eq. (40), whereas the two-halo contribution is likely to dominate on sufficiently large scales.

E. Galactic foregrounds

Foreground microwave emission from our own Galaxy can introduce spurious features in the weak lensing map. Because of their Galactic origin, these features cannot be correlated with the LRG distribution. However, it is possible that they can correlate with systematic errors in the LRG maps, most notably (i) stellar contamination of the LRG catalog and (ii) incomplete correction (or over-correction) for dust extinction. We have used two methods to address these potential problems. The first (Sec. VI E 1) is to correlate the CMB lensing map \mathbf{v} with stellar density and reddening maps. The second (Sec. VI E 2) is to correlate the LRG density map with simulated lensing contamination maps $\Delta\mathbf{v}$ obtained

TABLE VI: The bias obtained by substituting stellar density ($\delta n_\star/\bar{n}_\star$) and extinction [$E(B-V)$] maps in place of the LRG map ($\delta n_g/\bar{n}_g$), for the Kp05\\$\cap\\$S10\\$\backslash\\$ps2 sky cut. The final column, Δb_g , is the error in the galaxy bias if we use the coefficients c_E and c_E from the linear fit (Eq. 49) to estimate the contamination of the LRG map by stars and extinction.

Frequency	stars	$E(B-V)$	Δb_g
QQ	-14.38	-0.17	+0.02
QV	+2.37	+0.19	+0.10
QW	-2.12	+0.01	+0.03
VV	+3.40	+0.13	+0.05
VW	-2.29	+0.14	+0.11
WW	+3.99	+0.24	+0.11
TT	+0.71	+0.15	+0.09

by feeding microwave foreground templates through the lensing pipeline.

1. Stellar density and reddening tests

The dominant systematics in the SDSS that could correlate with Galactic microwave foregrounds are stellar contamination of the LRG catalog and dust extinction. We study these by constructing two maps: a map from SDSS of the density $\delta n_\star/\bar{n}_\star$ of “stars” (defined as objects with magnitude $18.0 < r < 19.5$ that are identified as pointlike by the SDSS photometric pipeline [51]), and a dust reddening map of $E(B-V)$ from Ref. [84]. These maps can be substituted in place of the LRG map $\delta n_g/\bar{n}_g$ as \mathbf{x}_{LRG} in Eq. (25) and the cross-spectrum and “bias” determined. The biases obtained using these contaminant maps are shown in Table VI.

A crude estimate of how this contamination translates into contamination of the galaxy-convergence power spectrum $C_l^{g\kappa}$ is provided by performing an unweighted least-squares fit of the LRG density to the reddening and stellar maps,

$$\frac{\delta n_g}{\bar{n}_g} = c_E E(B-V) + c_\star \frac{\delta n_\star}{\bar{n}_\star} + c_1 + \text{residual}, \quad (49)$$

over the 296,872 SDSS pixels. The fit coefficients are $c_E = 0.62$, $c_\star = -0.0092$, and $c_1 = -0.017$. We have shown in Table VI the spurious contribution Δb_g to the bias resulting from stellar and reddening contamination if one assumes these fit coefficients.

2. Microwave foreground template test

Since the lensing map \mathbf{v} is a quadratic function of temperature, and the Galactic foregrounds are not correlated with the primary CMB, Eq. (40) is applicable to Galactic foregrounds. The contamination $\Delta\langle\mathbf{v}^{\alpha\beta}\rangle$ can be obtained straightforwardly since the right-hand side of Eq. (40) can be evaluated by substituting in the foreground maps as ΔT^α . The difficult step is to construct

a good foreground map ΔT^α ; here we use external templates to avoid any possibility of spurious correlations of the templates with the *WMAP* data (either CMB signal or noise).

The Galactic foregrounds that must be considered in producing a template at higher (W band) frequencies are free-free and thermal dust emission; at lower frequencies (Q and V bands) an additional component is present whose physical origin remains uncertain but which may include hard synchrotron emission [83] or spinning or magnetic dust [99, 100]. We have used Model 8 of Refs. [84, 85] for thermal dust, and the $\text{H}\alpha$ line radiation template of Ref. [86] re-scaled using the conversions of Ref. [83] for free-free radiation. There are no all-sky synchrotron templates at the frequencies and angular scales of interest (the Haslam radio continuum maps at 408 MHz [87, 88], frequently used as a synchrotron template for CMB foreground analyses, have a 50 arcmin FWHM beam and hence do not resolve the $l \sim 200\text{--}400$ scales used for lensing of the CMB). Nevertheless, inclusion of the low-frequency component (whatever its origin) is not optional, and so we follow Ref. [100] in modeling it as proportional to the thermal dust prediction of Ref. [85] multiplied by T_{dust}^2 using the coefficients of Ref. [100].

As a test for contamination, we have substituted these foreground templates for the true CMB maps, run them through the lensing pipeline, and derived b_g estimates by correlating against the true LRG map; the results are shown in the “foreground” column of Table II. The typical contamination due to foregrounds is clearly very small (bias error of a few times 10^{-4}), and thus is negligible even if the foreground amplitude has been underestimated by an order of magnitude (the error on b_g scales as the foreground amplitude squared). This is not surprising: in the relatively clean regions of sky used for this analysis, the galactic foreground temperature anisotropy is roughly 2 orders of magnitude (in amplitude) below the CMB temperature at the degree angular scales. Therefore a quadratic statistic, such as the lensing estimator, should be ~ 4 orders of magnitude smaller than the foregrounds (again, in amplitude). Thus when $\mathbf{v}(\text{foreground})$ is correlated against the galaxy map, the correlation that one expects from chance alignments of foregrounds and galaxies is roughly 4 orders of magnitude less than $\mathbf{v}(\text{CMB})$ (although a much greater correlation could exist if the galaxy map were also contaminated by foregrounds, e.g. dust extinction). This is in contrast to the point sources, which are highly localized objects that become more and more dominant when we consider higher-order statistics such as the lensing estimator \mathbf{v} .

VII. DISCUSSION

In this paper, we have carried out an initial search for weak lensing of the CMB by performing a lensing reconstruction from the *WMAP* data and correlating the

resulting lensing field map with the SDSS LRG map. We do not have a detection, however our result $b_g = 1.81 \pm 1.92$ (1σ) is consistent with the bias $b_g \sim 1.8$ obtained from the LRG clustering autopower [60].

The main purpose of this analysis was to identify any systematics that contaminate the galaxy-convergence correlation at the level of the current CMB data. The good news is that our result for the bias is reasonable, suggesting that such systematics are at most of the order of the statistical errors. We have also found that the Galactic foregrounds are a negligible contaminant to the lensing signal (again, at the level of the present data). The bad news primarily concerns extragalactic foregrounds: a significant amount of solid angle – 21% of $\text{Kp05} \cap \text{S10}$ – was lost due to point source cuts that are necessary to avoid spurious power (at least in Q band), and we have no assurance that significant point source or tSZ contamination of the lensing signal does not lie just below the threshold of detectability. The extragalactic foreground analyses of Sec. VID 1 based on the frequency dependence of the signal and the galaxy-temperature correlations yielded only a weak constraint on the synchrotron point source contamination, $\Delta b_g^{(PS)}(TT) = 0.73 \pm 1.18$, and essentially no useful constraint can be derived for tSZ using the first-year *WMAP* data. Our constraints based on C_l^{gT} for the point sources are more stringent, $\Delta b_g^{(PS)} = -0.14 \pm 0.51$, but these assume Poissonianity of the sources, which must break down at some level. The point source and tSZ issues will become even more important as future experiments probe lensing of the CMB using higher- l primary modes, where point source and tSZ anisotropies contribute a greater fraction of the total power in the CMB, and precision cosmology with lensing of the CMB will require a means of constraining these contaminants in order to produce reliable results. Because in the real universe the extragalactic foregrounds will not be exactly Poisson-distributed, the frequency (in)dependence of the lensing signal will in principle provide the most robust constraints on the contamination. In this paper, we were unable to obtain useful constraints this way because of the limited range of frequencies on *WMAP* (all in the Rayleigh-Jeans regime where the tSZ signal has the same frequency dependence as CMB) and the low statistical signal-to-noise. Both of these problems should be alleviated with high-resolution data sets covering many frequencies, e.g. as expected from the *Planck* satellite [111].

The large solid angle that was lost to point source cuts in the analysis presented here resulted from the need to remove “artifacts” in the \mathbf{v} map that occur around point sources. One approach to this problem would be to try to devise a CMB lensing reconstruction technique that works with complicated masks. Alternatively, one could compute the galaxy density–CMB temperature–CMB temperature bispectrum $B_{l_1 l_2 l_3}^{gTT}$, rather than trying to use the lensing field (a quadratic function of CMB temperature) as an intermediate step; this way one could

mask out only the point source itself and not a 2 degree exclusion radius around it. The bispectrum approach carries the additional advantage of retaining angular information about the foregrounds; this information may be useful for separating lensing from the kinetic SZ and patchy-reionization anisotropies that have no frequency dependence but can still contaminate lensing if small-scale information is used [79, 94, 101]. The bispectrum $B_{l_1 l_2 l_3}^{gTT}$ may therefore be of particular interest for lensing analyses of high- l experiments such as the Atacama Pathfinder Experiment [112], the Atacama Cosmology Telescope [113], and the South Pole Telescope [114].

In summary, this paper represents a first analysis of lensing of the CMB using real data, and should not be regarded as the last word on the methodology. We have identified extragalactic foregrounds (point sources and tSZ) as the most worrying contaminant to the lensing signal in the *WMAP* data; the point sources, if unmasked, dominate the power spectrum of the reconstructed convergence if the Q band data are used, but this effect is suppressed at higher frequencies. We have shown that in the current data, the Galactic foreground contribution is negligible, and the contamination from point sources and tSZ in the galaxy-convergence cross-spectrum is at most of order the signal (although we have no detection of contamination). Like the search for the CMB lensing signal (and its eventual use in precision cosmology), stronger statements about the foreground contamination must await higher signal-to-noise data at a wide range of frequencies.

Acknowledgments

We acknowledge useful discussions with Niayesh Afshordi, Joseph Hennawi, Yeong-Shang Loh, and Lyman Page. C.H. is supported through NASA grant NGT5-50383. U.S. is supported by Packard Foundation, NASA NAG5-11489, and NSF CAREER-0132953.

Some of the results in this paper have been derived using the HEALPix [102] package. We acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA) [115]. Support for LAMBDA is provided by the NASA Office of Space Science. We used computational resources provided by NSF grant AST-0216105.

Funding for the creation and distribution of the SDSS Archive has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Aeronautics and Space Administration, the National Science Foundation, the U.S. Department of Energy, the Japanese Monbukagakusho, and the Max Planck Society. The SDSS Web site is <http://www.sdss.org/>.

The SDSS is managed by the Astrophysical Research Consortium (ARC) for the Participating Institutions. The Participating Institutions are The University of Chicago, Fermilab, the Institute for Advanced Study, the Japan Participation Group, The Johns Hop-

kins University, Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, University of Pittsburgh, Princeton University, the United States Naval Observatory, and the University of Washington.

APPENDIX A: NON-ISOLATITUDE SPHERICAL HARMONIC TRANSFORM

The non-isolatitude spherical harmonic transform (SHT) is used in our cross-correlation analysis. The SHT operations on the unit sphere transform between real- and harmonic-space representations of a function:

$$\begin{aligned} T(\hat{\mathbf{n}}_i) &= \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l T_{lm} Y_{lm}(\hat{\mathbf{n}}_i), & (\text{synthesis}) \\ S_{lm} &= \sum_{i=0}^{N_p-1} Y_{lm}^*(\hat{\mathbf{n}}_i) S(\hat{\mathbf{n}}_i). & (\text{analysis}) \end{aligned} \quad (A1)$$

For high-resolution data sets, this operation is usually performed using an isolatitude pixelization, i.e. one in which the pixels are positioned on curves of constant colatitude θ . This situation allows the colatitude (θ) and longitude (ϕ) parts of the spherical transform to be performed independently, resulting in an overall operation count scaling as $O(N_p^{3/2})$ [102, 103]. While this approach works, and has contributed remarkably to the popularity of isolatitude pixelizations such as HEALPix, there are reasons to maintain the flexibility to use any pixels. For example, in simulations of gravitational lensing of the CMB, we need to produce a simulated lensed map, and in general a set of pixels that are isolatitude in “observed” coordinates (e.g. HEALPix) maps onto a non-isolatitude grid on the primary CMB. We note that for other analyses there may be other reasons to consider more general pixelizations, which preserve desired properties such as conformality [104] or maximal symmetry [105]. This Appendix describes our non-isolatitude SHT algorithm.

1. The method

We consider first the SHT synthesis. Our first step is to perform a latitude transform using associated Legendre polynomials on a set of points equally spaced in θ (the “coarse grid”):

$$T_m\left(\frac{\alpha}{L}\pi\right) = \sum_{l=|m|}^{l_{\max}} T_{lm} Y_{lm}\left(\theta = \frac{\alpha}{L}\pi, \phi = 0\right). \quad (A2)$$

This procedure is performed for integers in the range $0 \leq \alpha \leq L$, $-l_{\max} \leq m \leq l_{\max}$. [Here L is an integer satisfying $L > l_{\max}$, which we require to be a power of two times a small odd integer. The first requirement ensures that Eq. (A2) over-Nyquist samples the variations in $T_m(\theta)$, the second ensures that the FFT is a fast operation.] It requires a total of $O(l_{\max}^2 L)$ operations and

is the most computationally demanding step in the transform. The spherical harmonics are computed as needed using an ascending recursion relation. The standard recursion relation is used to generate the associated Legendre functions; we speed up the transform by a factor of two over the brute-force approach by taking advantage of the symmetry/antisymmetry of the spherical harmonics across the equatorial plane.

The next step is to refine the coarse grid, which has a spacing of π/L in θ , to a “fine grid” with spacing π/L' where $L' > L$. We do this by taking advantage of the band-limited nature of the spherical harmonics. Any linear combination of spherical harmonics of order $l \leq l_{\max}$ can be written as a band-limited function:

$$T_m(\theta) = \sum_{l=|m|}^{l_{\max}} T_{lm} Y_{lm}(\theta, 0) = \sum_{n=-l_{\max}}^{l_{\max}} C_{m,n} e^{in\theta}. \quad (\text{A3})$$

We may determine the coefficients $C_{m,n}$ via a fast Fourier transform (FFT) of length $2L$, so long as the left-hand side has been evaluated at the points $\theta = \pi\alpha/L$ for integers $-L < \alpha \leq L$. (We use parity rules to compute the left hand side at negative values of θ .) By applying an FFT of length $2L'$ to the $C_{m,n}$, we then obtain $T_m(\theta)$ at values of $\theta = \pi\alpha/L'$ for integers $0 \leq \alpha \leq L'$. What has been gained here is that we have performed the associated Legendre transform on $2L'$ points, but the expensive evaluation of the associated Legendre polynomials has only been required at $2L$ points. If L and L' are powers of two, then the FFT process requires $O(l_{\max}L' \log L')$ operations.

The third step is an FFT in the longitude direction to obtain $T(\theta = \pi\alpha/L', \phi = \pi\gamma/L')$, where α and γ are integers. This process has become standard in isolatitude SHT algorithms; in its full glory, it is given by:

$$T[\alpha, \gamma] \equiv T(\theta = \frac{\alpha}{L'}\pi, \phi = \frac{\gamma}{L'}\pi) = \sum_{m=-l_{\max}}^{l_{\max}} T_m(\theta) e^{im\phi}. \quad (\text{A4})$$

At the end of this step, we know the real-space value of our function on a fine equicylindrical projection (ECP) grid of spacing π/L' . The FFT operation in this step requires $O(L'^2 \log L')$ operations. The total operation count of transforming onto the ECP grid is $O(l_{\max}^3)$ and is dominated by the associated Legendre transform, so long as $L'/l_{\max} < \sqrt{l_{\max}}/\log l_{\max}$. What is important to note is that, particularly if l_{\max} is large (for a future CMB polarization experiment we have to consider multipoles up to roughly $l_{\max} \sim 4000$), we can sample our function in real space at several times the Nyquist frequency at no additional computational cost. This is exactly what is required in order to successfully interpolate the values T_j .

The final step is the interpolation step. For each point $\hat{\mathbf{n}}_j$, we identify the coordinates in the ECP grid: (we suppress the j index here for clarity)

$$\alpha + \delta_\alpha = L' \frac{\theta}{\pi}; \gamma + \delta_\gamma = L' \frac{\phi}{\pi}, \quad (\text{A5})$$

where α and γ are integers and the fractional parts $0 \leq \delta_\alpha, \delta_\gamma < 1$. A $4K^2$ -point, two-dimensional polynomial interpolation is then computed:

$$T \approx \sum_{\mu=-K+1}^K w_\mu(\delta_\alpha) \sum_{\nu=-K+1}^K w_\nu(\delta_\gamma) T[\alpha+\mu, \gamma+\nu], \quad (\text{A6})$$

where the weights $w_\rho(\delta)$ are computed by Lagrange’s formula:

$$w_\rho(\delta) = \frac{(-1)^{K-\rho}}{(K-\rho)!(K-1+\rho)!(\delta-\rho)} \prod_{\sigma=-K+1}^K (\delta - \sigma). \quad (\text{A7})$$

The weights for both the α and γ directions may be evaluated in a total of $O(K)$ multiplications and divisions if the factorials have been pre-computed, so that for high-order interpolations the dominant contribution to the computation time in interpolation comes from the multiplications in Eq. (A6) rather than from computation of the weights.

Note that the “analysis” operation of Eq. (A1) is the matrix transpose operation of the “synthesis” (if we view the pixelized $T(\hat{\mathbf{n}}_i)$ and the harmonic-space T_{lm} as vectors), and that all of the steps outlined above are linear operations. Since the transpose of a composition of operations is the composition of the transposes in reverse order, we can simply use the transposes of these operations in reverse order to compute an SHT analysis using the same number of operations.

The above SHT algorithm generalizes easily to vector spherical harmonics, for which we use the basis:

$$Y_{lm}^{(\parallel)}(\hat{\mathbf{n}}) = \frac{1}{\sqrt{l(l+1)}} \nabla Y_{lm}(\hat{\mathbf{n}}),$$

$$Y_{lm}^{(\perp)}(\hat{\mathbf{n}}) = \frac{1}{\sqrt{l(l+1)}} \hat{\mathbf{n}} \times \nabla Y_{lm}(\hat{\mathbf{n}}). \quad (\text{A8})$$

2. Interpolation accuracy

The order K and step size π/L' of the interpolating polynomial is determined by a balance of computation time, memory usage, and accuracy. The computational cost of the interpolation is $O(K^2 N_{\text{pix}})$ and the memory usage is $O(L'^2)$. Hence it is important to understand how interpolation accuracy is related to K and L' .

The error in polynomial interpolation of $T(\theta, \phi)$ can be determined from the error in interpolation of a band-limited function. We note that the error in interpolation of a given Fourier mode $f(\theta, \phi) = C_{m,n} e^{i(n\theta+m\phi)}$ is given by:

$$\frac{f_{\text{interp}}(\theta, \phi)}{f(\theta, \phi)} = \left[1 + v\left(\delta_\alpha, 2\pi \frac{m}{L'}\right) \right] \left[1 + v\left(\delta_\gamma, 2\pi \frac{n}{L'}\right) \right], \quad (\text{A9})$$

where the v function is:

$$v(\delta, \psi) = \sum_{\rho=-K+1}^K w_\rho(\delta) e^{i\psi(\rho-\delta)} - 1. \quad (\text{A10})$$

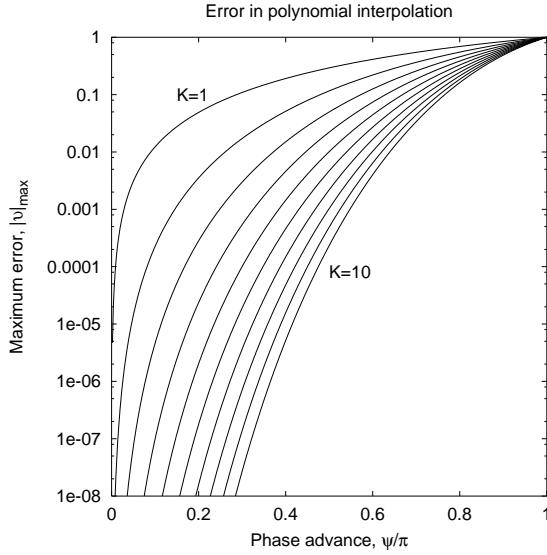


FIG. 13: The maximum fractional error $|v_{\max}|$ from Eq. (A12) as a function of the interpolation order K and the sampling rate. From top to bottom, the curves correspond to $K = 1$ to $K = 10$. Note that $\psi = \pi$ for data sampled at the Nyquist frequency (i.e. $L' = l_{\max}$), whereas $\psi = \pi/2$ for data sampled at twice the Nyquist frequency ($L' = 2l_{\max}$), etc. Accuracy can be improved via either fine sampling (small ψ) or high-order interpolation (high K).

(Note that ψ represents the phase advance of our Fourier mode per grid cell.) The v function determines the fractional error in a given Fourier mode. It is most easily evaluated by noting that (since polynomial interpolation is exact for a constant) we have $v(\delta, 0) = 0$, and derivative:

$$\left| \frac{\partial v}{\partial \psi} \right| = \frac{\prod_{\sigma=-K+1}^K |\delta - \sigma|}{(2K-1)!} |1 - e^{-i\psi}|^{2K-1}. \quad (\text{A11})$$

[Here we have used Eq. (A7) and applied the binomial theorem.] We note that the product is maximized at $\delta = 1/2$ and that using trigonometric identities the exponential term can be simplified. Since $v(\delta, 0) = 0$, we have that $|v|$ cannot exceed the integral of $|dv/d\psi|$ from 0 to ψ :

$$|v(\delta, \psi)| \leq \frac{(2K)!}{4^K K! (K-1)!} \int_0^\psi \sin^{2K-1} \frac{\psi'}{2} d\psi'. \quad (\text{A12})$$

This bound is plotted in Fig. 13.

APPENDIX B: PRECONDITIONER

One of the steps in the cross-spectrum estimator requires the solution of the linear system:

$$\mathbf{C}_p \mathbf{y} = \mathbf{x} \quad (\text{B1})$$

for \mathbf{y} . Here \mathbf{C}_p is the prior covariance matrix for the LRGs and is given by $\mathbf{C}_p = \mathbf{S}_p + \mathbf{N}$, where \mathbf{S}_p is the signal

prior and \mathbf{N} is the noise. Since the matrix \mathbf{C}_p is too large ($N(N+1)/2$ elements where the dimension $N = 296,872$) to store in memory, Eq. (B1) must be solved using iterative methods. We have used a preconditioned conjugate-gradient (PCG) method. The PCG method [106, 107] with preconditioner \mathbf{E} produces successive estimates $\mathbf{y}^{(i)}$ for the solution to Eq. (B1) using the equation:

$$\mathbf{y}^{(i)} = \mathbf{y}^{(i-1)} + \frac{\mathbf{r}^{(i-1)T} \mathbf{E} \mathbf{r}^{(i-1)}}{\mathbf{p}^{(i)} \mathbf{C}_p \mathbf{p}^{(i)}} \mathbf{p}^{(i)}, \quad (\text{B2})$$

where the residuals are defined by $\mathbf{r}^{(i)} = \mathbf{x} - \mathbf{C}_p \mathbf{y}^{(i)}$, and the search directions $\mathbf{p}^{(i)}$ are chosen according to

$$\mathbf{p}^{(i)} = \mathbf{E} \mathbf{r}^{(i-1)} + \frac{\mathbf{r}^{(i-1)T} \mathbf{E} \mathbf{r}^{(i-1)}}{\mathbf{r}^{(i-2)T} \mathbf{E} \mathbf{r}^{(i-2)}} \mathbf{p}^{(i-1)}. \quad (\text{B3})$$

The initial conditions are $\mathbf{y}^{(0)} = 0$ and $\mathbf{p}^{(1)} = \mathbf{E} \mathbf{x}$.

The choice of preconditioner strongly affects the rate of convergence of the PCG algorithm; ideally we have a preconditioner for which $\mathbf{E} \mathbf{u}$ can be rapidly computed for given \mathbf{u} , and for which $\mathbf{E} \approx \mathbf{C}_p^{-1}$. In using PCG for CMB power spectrum estimation, Ref. [106] used a preconditioner based on (approximate) azimuthal symmetry of the Galactic plane cut. Unfortunately, the SDSS survey region lacks any such symmetry. Therefore we use a preconditioner that works independently of any symmetries of the survey region. The strategy is to use a two-scale preconditioner: the low-resolution scale ($l < l_{\text{split}}$) is solved by brute-force matrix inversion in harmonic space, while the high-resolution scale ($l \geq l_{\text{split}}$) is essentially un-preconditioned and the burden of convergence lies on the conjugate-gradient algorithm. This is useful because the large condition ratio of a typical prior matrix \mathbf{C}_p comes almost entirely from a few large eigenvalues corresponding to the low l modes. The condition ratio is dramatically improved by suppressing these eigenvalues with the two-scale preconditioner.

The preconditioner is obtained by cutting the power spectrum at l_{split} and defining the $l_{\text{split}}^2 \times l_{\text{split}}^2$ matrix \mathbf{M} :

$$\begin{aligned} M_{lm, l'm'} &= \delta_{ll'} \delta_{mm'} \\ &+ \frac{\sqrt{(S_{p,l} - S_{p,l_{\text{split}}})(S_{p,l'} - S_{p,l_{\text{split}}})}}{S_{p,l_{\text{split}}}} \\ &\times \sum_j \Omega_j Y_{lm}^*(\hat{\mathbf{n}}_j) Y_{l'm'}(\hat{\mathbf{n}}_j). \end{aligned} \quad (\text{B4})$$

Here l and l' are in the range from 0 to $l_{\text{split}} - 1$, j is a pixel index, and Ω_j is the area of the j th pixel. Note that a red prior spectrum (i.e. $S_{p,l} > S_{p,l_{\text{split}}}$) has been assumed, as appropriate for the angular power spectra of galaxies. We next define the \mathbf{G} matrix:

$$\begin{aligned} G_{lm, l'm'} &= \frac{\sqrt{(S_{p,l} - S_{p,l_{\text{split}}})(S_{p,l'} - S_{p,l_{\text{split}}})}}{S_{p,l_{\text{split}}}} \\ &\times [\mathbf{M}^{-1}]_{lm, l'm'}. \end{aligned} \quad (\text{B5})$$

Then the preconditioner is:

$$E_{ij} = \frac{\Omega_i \delta_{ij}}{S_{p,l_{\text{split}}}} - \sum_{lml'm'} \frac{\Omega_i \Omega_j}{S_{p,l_{\text{split}}}^2} \times Y_{lm}(\hat{\mathbf{n}}_i) Y_{l'm'}^*(\hat{\mathbf{n}}_j) G_{lm,l'm'}. \quad (\text{B6})$$

Using the Woodbury matrix inverse formula, the matrix \mathbf{E} of Eq. (B6) can be recognized as the exact inverse of the matrix:

$$[\mathbf{E}^{-1}]_{ij} = S_{p,l_{\text{split}}} \frac{\delta_{ij}}{\Omega_i} + \sum_{l=0}^{l_{\text{split}}-1} (S_{p,l} - S_{p,l_{\text{split}}}) \times \sum_{m=-l}^l Y_{lm}(\hat{\mathbf{n}}_i) Y_{lm}^*(\hat{\mathbf{n}}_j). \quad (\text{B7})$$

That is, \mathbf{E} is the exact inverse of a convolution that has power spectrum $S_{p,l}$ for $l < l_{\text{split}}$ and white noise with power $S_{p,l_{\text{split}}}$ for $l \geq l_{\text{split}}$. Thus \mathbf{E} is a good approximation to \mathbf{C}_p if l_{split} is chosen to be the multipole where the signal and noise have similar power. In this case,

the convergence of the PCG iteration is very rapid. In practice (see below) we choose l_{split} to be somewhat less than this in order to speed up multiplication of a vector by \mathbf{E} , and accept that more iterations will be required for convergence.

Note that \mathbf{L} need only be computed once for each sky cut and prior power spectrum; it can then be stored and used for many $\mathbf{E}\mathbf{u}$ operations. The $\mathbf{E}\mathbf{u}$ operation is then reduced to a spherical harmonic transform of cost $O(l_{\text{split}}^3) + O(K^2 N_{\text{pix}})$, a matrix-vector multiplication of computational cost $O(l_{\text{split}}^4)$, and another spherical harmonic transform. Since we use the preconditioner once for every \mathbf{C}_p operation, we gain speed by increasing l_{split} until the most expensive part of the \mathbf{E} operation is of comparable cost to the \mathbf{C}_p operation ($O(l_{\text{max}}^3)$). This suggests that we set $l_{\text{split}} \sim l_{\text{max}}^{3/4}$. There are however practical limitations on l_{split} : the size of the matrix in memory is $O(l_{\text{split}}^4)$, and the computational cost of obtaining \mathbf{G} is $O(l_{\text{split}}^6)$. Thus the best value of l_{split} is generally somewhat less than $l_{\text{max}}^{3/4}$; we have used $l_{\text{split}} = 32$.

[1] C. L. Bennett et al., *Astrophys. J. Supp.* **148**, 1 (2003).
[2] G. Hinshaw et al., *Astrophys. J. Supp.* **148**, 135 (2003).
[3] A. Kogut et al., *Astrophys. J. Supp.* **148**, 161 (2003).
[4] D. N. Spergel et al., *Astrophys. J. Supp.* **148**, 175 (2003).
[5] L. Verde et al., *Astrophys. J. Supp.* **148**, 195 (2003).
[6] H. V. Peiris et al., *Astrophys. J. Supp.* **148**, 213 (2003).
[7] L. Page et al., *Astrophys. J. Supp.* **148**, 233 (2003).
[8] P. Fosalba, E. Gaztañaga, and F. J. Castander, *Astrophys. J. Lett.* **597**, L89 (2003).
[9] J. M. Diego, J. Silk, and W. Sliwa, *Mon. Not. R. Astron. Soc.* **346**, 940 (2003).
[10] S. Boughn and R. Crittenden, *Nature (London)* **427**, 45 (2004).
[11] A. D. Myers, T. Shanks, P. J. Outram, W. J. Frith, and A. W. Wolfendale, *Mon. Not. R. Astron. Soc.* **347**, L67 (2004).
[12] M. R. Nolta et al., *Astrophys. J.* **608**, 10 (2004).
[13] R. Scranton et al. (2003), astro-ph/0307335.
[14] N. Afshordi, Y. Loh, and M. A. Strauss, *Phys. Rev. D* **69**, 083524 (2004).
[15] A. Refregier, *Annu. Rev. Astron. Astrophys.* **41**, 645 (2003).
[16] L. Van Waerbeke et al., *Astron. Astrophys.* **358**, 30 (2000).
[17] D. J. Bacon, A. R. Refregier, and R. S. Ellis, *Mon. Not. R. Astron. Soc.* **318**, 625 (2000).
[18] J. Rhodes, A. Refregier, and E. J. Groth, *Astrophys. J. Lett.* **552**, L85 (2001).
[19] H. Hoekstra, H. K. C. Yee, M. D. Gladders, L. F. Barrientos, P. B. Hall, and L. Infante, *Astrophys. J.* **572**, 55 (2002).
[20] L. Van Waerbeke, Y. Mellier, R. Pelló, U.-L. Pen, H. J. McCracken, and B. Jain, *Astron. Astrophys.* **393**, 369 (2002).
[21] M. Jarvis, G. M. Bernstein, P. Fischer, D. Smith, B. Jain, J. A. Tyson, and D. Wittman, *Astron. J.* **125**, 1014 (2003).
[22] M. L. Brown, A. N. Taylor, D. J. Bacon, M. E. Gray, S. Dye, K. Meisenheimer, and C. Wolf, *Mon. Not. R. Astron. Soc.* **341**, 100 (2003).
[23] R. Massey, A. Refregier, D. Bacon, and R. Ellis (2004), astro-ph/0404195.
[24] T. G. Braine, R. D. Blandford, and I. Smail, *Astrophys. J.* **466**, 623 (1996).
[25] H. Hoekstra, L. van Waerbeke, M. D. Gladders, Y. Mellier, and H. K. C. Yee, *Astrophys. J.* **577**, 604 (2002).
[26] P. Fischer et al., *Astron. J.* **120**, 1198 (2000).
[27] T. A. McKay et al. (2001), astro-ph/0108013.
[28] E. S. Sheldon et al., *Astron. J.* **127**, 2544 (2004).
[29] H. Hoekstra, H. K. C. Yee, and M. D. Gladders, *Astrophys. J.* **606**, 67 (2004).
[30] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
[31] W. Hu, *Phys. Rev. D* **65**, 023003 (2002).
[32] M. Kaplinghat, *New Astronomy Review* **47**, 893 (2003).
[33] W. Hu and T. Okamoto, *Phys. Rev. D* **69**, 043004 (2004).
[34] A. Lewis, *Phys. Rev. D* **70**, 043518 (2004).
[35] L. Knox and Y. Song, *Physical Review Letters* **89**, 11303 (2002).
[36] M. Kesden, A. Cooray, and M. Kamionkowski, *Physical Review Letters* **89**, 11304 (2002).
[37] U. Seljak and C. M. Hirata, *Phys. Rev. D* **69**, 043005 (2004).
[38] C. L. Bennett et al., *Astrophys. J.* **583**, 1 (2003).
[39] N. Jarosik et al., *Astrophys. J. Supp.* **148**, 29 (2003).
[40] G. Hinshaw et al., *Astrophys. J. Supp.* **148**, 63 (2003).
[41] M. Tegmark, A. de Oliveira-Costa, and A. J. Hamilton, *Phys. Rev. D* **68**, 123523 (2003).
[42] D. G. York et al., *Astron. J.* **120**, 1579 (2000).
[43] M. A. Strauss et al., *Astron. J.* **124**, 1810 (2002).
[44] G. T. Richards et al., *Astron. J.* **123**, 2945 (2002).

[45] D. W. Hogg, D. P. Finkbeiner, D. J. Schlegel, and J. E. Gunn, *Astron. J.* **122**, 2129 (2001).

[46] M. Fukugita, T. Ichikawa, J. E. Gunn, M. Doi, K. Shimasaku, and D. P. Schneider, *Astron. J.* **111**, 1748 (1996).

[47] J. A. Smith et al., *Astron. J.* **123**, 2121 (2002).

[48] J. E. Gunn et al., *Astron. J.* **116**, 3040 (1998).

[49] M. R. Blanton, H. Lin, R. H. Lupton, F. M. Maley, N. Young, I. Zehavi, and J. Loveday, *Astron. J.* **125**, 2276 (2003).

[50] J. R. Pier, J. A. Munn, R. B. Hindsley, G. S. Hennessy, S. M. Kent, R. H. Lupton, and Ž. Ivezić, *Astron. J.* **125**, 1559 (2003).

[51] R. H. Lupton, J. E. Gunn, Ž. Ivezić, G. R. Knapp, S. Kent, and N. Yasuda, in *ASP Conf. Ser. 238: Astronomical Data Analysis Software and Systems X* (2001), pp. 269–278.

[52] C. Stoughton et al., *Astron. J.* **123**, 485 (2002).

[53] K. Abazajian et al., *Astron. J.* **126**, 2081 (2003).

[54] K. Abazajian et al., *Astron. J.* **128**, 502 (2004).

[55] D. Finkbeiner et al. (2004), *Astron. J.*, accepted.

[56] H. V. Peiris and D. N. Spergel, *Astrophys. J.* **540**, 605 (2000).

[57] M. D. Gladders and H. K. C. Yee, *Astron. J.* **120**, 2148 (2000).

[58] D. J. Eisenstein et al., *Astron. J.* **122**, 2267 (2001).

[59] N. Padmanabhan, et al. (2004), astro-ph/0407594.

[60] N. Padmanabhan et al. (2004), in preparation.

[61] E. Høg, C. Fabricius, V. V. Makarov, S. Urban, T. Corbin, G. Wycoff, U. Bastian, P. Schwekendiek, and A. Wicenec, *Astron. Astrophys.* **355**, L27 (2000).

[62] W. Hu, *Phys. Rev. D* **64**, 083005 (2001).

[63] T. Okamoto and W. Hu, *Phys. Rev. D* **67**, 83002 (2003).

[64] L. Page et al., *Astrophys. J. Supp.* **148**, 39 (2003).

[65] U. Seljak and M. Zaldarriaga, *Astrophys. J.* **469**, 437 (1996).

[66] M. Tegmark et al., *Phys. Rev. D* **69**, 103501 (2004).

[67] J. A. Peacock and S. J. Dodds, *Mon. Not. R. Astron. Soc.* **280**, L19 (1996).

[68] J. N. Fry, *Astrophys. J. Lett.* **461**, L65+ (1996).

[69] F. Bernardeau, *Astron. Astrophys.* **324**, 15 (1997).

[70] U. Seljak and M. Zaldarriaga, *Physical Review Letters* **82**, 2636 (1999).

[71] J. Guzik, U. Seljak, and M. Zaldarriaga, *Phys. Rev. D* **62**, 45317 (2000).

[72] M. Zaldarriaga, *Phys. Rev. D* **62**, 63510 (2000).

[73] W. Hu, *Astrophys. J.* **557**, L79 (2001).

[74] T. Okamoto and W. Hu, *Phys. Rev. D* **66**, 063008 (2002).

[75] A. Cooray, *New Astronomy* **9**, 173 (2004).

[76] C. M. Hirata and U. Seljak, *Phys. Rev. D* **67**, 43001 (2003).

[77] C. M. Hirata and U. Seljak, *Phys. Rev. D* **68**, 83002 (2003).

[78] U. Seljak and M. Zaldarriaga, *Astrophys. J.* **538**, 57 (2000).

[79] C. Vale, A. Amblard, and M. White (2004), astro-ph/0402004.

[80] S. Dodelson, *Phys. Rev. D* **70**, 023009 (2004).

[81] C.-L. Kuo et al., American Astronomical Society Meeting **201** (2002).

[82] T. J. Pearson and Cosmic Background Imager Collaboration, American Astronomical Society Meeting **200** (2002).

[83] C. L. Bennett et al., *Astrophys. J. Supp.* **148**, 97 (2003).

[84] D. J. Schlegel, D. P. Finkbeiner, and M. Davis, *Astrophys. J.* **500**, 525 (1998).

[85] D. P. Finkbeiner, M. Davis, and D. J. Schlegel, *Astrophys. J.* **524**, 867 (1999).

[86] D. P. Finkbeiner, *Astrophys. J. Supp.* **146**, 407 (2003).

[87] C. G. T. Haslam, U. Klein, C. J. Salter, H. Stoffel, W. E. Wilson, M. N. Cleary, D. J. Cooke, and P. Thomasson, *Astron. Astrophys.* **100**, 209 (1981).

[88] C. G. T. Haslam, H. Stoffel, C. J. Salter, and W. E. Wilson, *Astron. Astrophys. Supp.* **47**, 1 (1982).

[89] A. J. S. Hamilton, *Mon. Not. R. Astron. Soc.* **289**, 285 (1997).

[90] A. J. S. Hamilton, *Mon. Not. R. Astron. Soc.* **289**, 295 (1997).

[91] M. Tegmark, *Phys. Rev. D* **55**, 5895 (1997).

[92] N. Padmanabhan, U. Seljak, and U. L. Pen, *New Astronomy* **8**, 581 (2003).

[93] C. M. Hirata et al. (2004), astro-ph/0403255.

[94] A. Amblard, C. Vale, and M. White (2004), astro-ph/0403075.

[95] A. Cooray and M. Kesden, *New Astronomy* **8**, 231 (2003).

[96] M. Kesden, A. Cooray, and M. Kamionkowski, *Phys. Rev. D* **67**, 123507 (2003).

[97] E. Komatsu et al., *Astrophys. J. Supp.* **148**, 119 (2003).

[98] N. Padmanabhan et al. (2004), in preparation.

[99] A. de Oliveira-Costa, M. Tegmark, R. D. Davies, C. M. Gutiérrez, A. N. Lasenby, R. Rebolo, and R. A. Watson, *Astrophys. J. Lett.* **606**, L89 (2004).

[100] D. P. Finkbeiner (2003), astro-ph/0311547.

[101] M. G. Santos, A. Cooray, Z. Haiman, L. Knox, and C. Ma, *Astrophys. J.* **598**, 756 (2003).

[102] K. M. Górski and et al., in *Evolution of Large Scale Structure : From Recombination to Garching* (1999), pp. 37–+.

[103] R. G. Crittenden and N. G. Turok (1998), astro-ph/9806374.

[104] W. N. Colley and J. R. Gott III (2003), astro-ph/0303020.

[105] M. Tegmark, *Astrophys. J. Lett.* **470**, L81 (1996).

[106] S. P. Oh, D. N. Spergel, and G. Hinshaw, *Astrophys. J.* **510**, 551 (1999).

[107] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C. The art of scientific computing* (Cambridge: University Press, —c1992, 2nd ed., 1992).

[108] URL: <http://map.gsfc.nasa.gov/>

[109] URL: <http://www.sdss.org/>

[110] URL: <http://www.eso.org/science/healpix/>

[111] URL: <http://astro.estec.esa.nl/Planck>

[112] URL: <http://bolo.berkeley.edu/apexsz/>

[113] URL: <http://www.hep.upenn.edu/~angelica/act/act.html>

[114] URL: <http://astro.uchicago.edu/spt/>

[115] URL: <http://lambda.gsfc.nasa.gov/>